ESSAYS ON DEBT, FINANCIAL CRISIS, AND IMPULSE RESPONSE FUNCTIONS

by

Yun Jung Kim

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Economics)
in The University of Michigan
2012

Doctoral Committee:

Professor Linda L. Tesar, Co-Chair
Assistant Professor Jing Zhang, Co-Chair
Professor Kathryn Mary Dominguez
Associate Professor Uday Rajan

Table of Contents

List of Figures

List of Tables

CHAPTER I

Decentralized Borrowing and Centralized Default

with Jing Zhang

1.1 Introduction

In the past, foreign borrowing by developing countries was comprised almost entirely of government borrowing. Motivated by this observation, the sovereign debt literature focuses on "*centralized borrowing*" models in which governments decide both how much to borrow and whether to repay. In recent decades private external borrowing has risen substantially from less than 20 percent of total external borrowing in 1990 to more than 70 percent in 2008, as shown in Figure 1. Private external debt is often priced with macroeconomic indicators rather than according to individual borrowers' ability to repay because governments play an important role in private external debt repayments.[1] In such an environment, a pecuniary externality arises because private agents fail to internalize the impacts of their individual borrowing on credit costs. Intuitively, such a pecuniary externality will lead to excessive borrowing.

This paper quantitatively evaluates the impact of this pecuniary externality in a *decentralized borrowing* model where private agents decide how much to borrow, and

[1] One reason that governments affect repayments of private external debt is that developing countries' external debt is predominantly denominated in foreign currencies and governments control exchange rates. A devaluation of domestic currencies can cause wide-spread defaults on private external debt. Another reason is that governments often explicitly or implicitly guarantee private external debt. Chile's debt nationalization in 1982 is an example. Chilean total foreign debt reached as high as 20 billion dollars in 1982, and two thirds was private debt by leading domestic private banks. When the Latin American panic dried up new loans, six top private banks failed, and the Chilean government assumed responsibility for private foreign debt. As a result, when pricing private external loans, foreign lenders rely heavily on macroeconomic indicators of developing countries. See Section 1.2.1 for more discussions.

Figure 1: Debt Flows to Private Sectors of Developing Countries

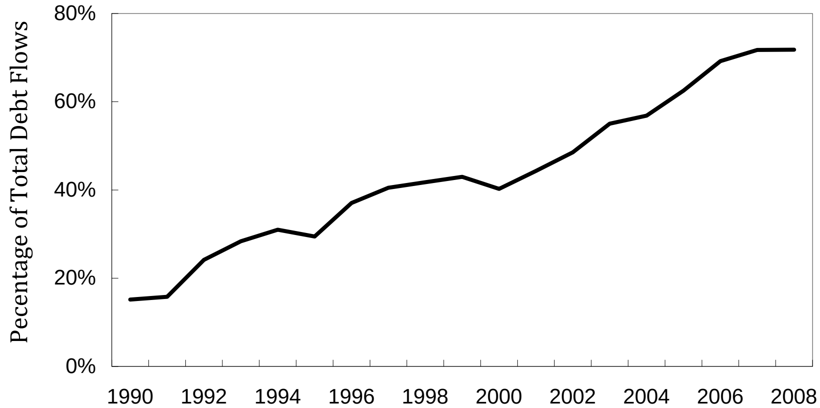

Data source: *Global Development Finance*

the government decides whether to default. Interestingly, the model suggests that the quantitative effect of this pecuniary externality on debt levels is modest and ambiguous. Equilibrium debt levels could be higher or lower under decentralized borrowing than under centralized borrowing, depending on the specification and severity of default penalties. On the other hand, decentralized borrowing unambiguously generates higher credit costs, larger default frequencies and lower welfare.

In our model, a continuum of identical households borrow non-contingent debt from foreign lenders to smooth their income shocks. A benevolent government decides whether to enforce debt repayments to maximize the welfare of the representative household. If the government defaults, the country loses access to international financial markets and suffers income losses for some stochastic number of periods. Foreign lenders offer the households a price schedule for debt that depends on the level of aggregate borrowing instead of individual borrowing. This is because aggregate borrowing, together with the aggregate income shock, determines the default probability. We calibrate the model to Argentine data to examine quantitative effects of decentralized borrowing.

The paper identifies two channels through which decentralized borrowing affects debt decisions relative to the centralized borrowing model. The first is straightforward: private agents do not internalize the adverse effect of an extra unit of debt on government default probabilities and aggregate credit costs, and therefore borrow more for a given interest rate schedule. This *overborrowing effect* shifts up the demand curve for debt and tends to increase equilibrium debt levels as well as credit costs. The second channel is less obvious: the equilibrium interest rate schedule is higher under decentralized borrowing, inducing private agents to borrow less. The interest rate schedule rises because the anticipation of future overborrowing by private agents leads to larger default likelihood of the government for any debt level. This *bond price schedule effect* shifts up the supply curve of debt and tends to reduce debt levels while at the same time increasing credit costs. Thus, equilibrium debt levels depend on the relative strength of the two channels.

Since default penalties play a central role in quantitative results, we examine two commonly-used specifications of default penalties. The first is an *asymmetric default penalty*, in which income losses after default are disproportionately large under good income shocks.[2] The second is a *symmetric default penalty*, in which income losses after default are a constant share of income. Under asymmetric default penalties, decentralized borrowing generates higher debt levels than centralized borrowing when default penalties are lenient, but lower debt levels when default penalties are harsh. Under symmetric default penalties, equilibrium debt levels are always lower under decentralized borrowing.

To understand the impact of decentralized borrowing on equilibrium debt levels, let us consider the two counteracting effects determining equilibrium debt. The overborrowing effect arises from the failure of private agents to internalize the impact of their additional borrowing on credit costs, so operates only when agents take on risky debt. The more sensitive credit costs are to aggregate borrowing, the stronger is the

[2]Proposed by Arellano (2008), this specification of default penalties helps the centralized borrowing model generate an empirically reasonable default rate.

overborrowing effect. The bond price schedule effect occurs because the overborrowing incentive lowers the repayment welfare and drives up the default likelihood of the government. The greater the decrease in the repayment welfare is, the stronger is the bond price schedule effect.

With symmetric default penalties, the overborrowing effect barely operates because agents rarely borrow risky debt. Thus, the bond price schedule effect dominates, leading to equilibrium underborrowing. With asymmetric default penalties, agents are more likely to borrow risky debt and the overborrowing effect operates. When default penalties are lenient, credit costs are sensitive to aggregate risky borrowing because the government's incentive to default rises rapidly with additional debt. The reduction in the repayment welfare from decentralized borrowing is small because default is not that costly. Thus, the overborrowing effect is strong while the bond price schedule effect is weak, which leads to equilibrium overborrowing. When default penalties are severe, credit costs are less sensitive to risky borrowing, but the repayment welfare is lowered substantially under decentralized borrowing. Thus, the overborrowing effect is weak while the bond price schedule effect is strong, which leads to equilibrium underborrowing.

Our work is related to Jeske (2006) and Wright (2006), who study theoretically the impact of decentralized borrowing in an environment with complete markets and default risk. Our paper examines such effects quantitatively in an environment with incomplete markets and default risk.[3] Our work relates to many studies that analyze the effect of pecuniary externalities coming from other sources in debt markets. Bizer and DeMarzo (1992) studies an externality arising from sequential borrowing from multiple lenders. Bi (2006) and Hatchondo and Martinez (2009) examine the Bizer-DeMarzo type of externality in quantitative sovereign debt models. Lorenzoni (2008) studies an externality arising from failure of private investors to take into account the

[3]Bai and Zhang (2010) show that both incomplete markets and default risk are important to account for various dimensions of international data, for example, savings and investment behavior.

effect of private asset sales on asset prices.

Our work is also related to Uribe (2006) who shows that regardless of whether a debt limit is imposed at the country level or at the individual level, the equilibrium level of debt is the same. In his analysis, debt limits and interest rates are exogenously specified, and there is no default risk. By contrast, in our model, the presence of default risk endogenizes both interest rates and debt limits, and whether interest rates depend on aggregate debt or individual debt affects the equilibrium level of debt.[4]

Our model builds on the classic sovereign default framework of Eaton and Gersovitz (1981) and recent quantitative research on sovereign debt: Arellano (2008), Aguiar and Gopinath (2006), among others. Recently, different approaches have been taken to enrich and improve the sovereign debt model. Bai and Zhang (2009)) introduce production economies to the sovereign debt literature. Cuadra and Sapriza (2008) and Hatchondo et al. (2009) incorporate political economy considerations into the government's decision. Arellano and Ramanarayan (2010), Chatterjee and Eyigungor (2010) and Hatchondo and Martinez (2009) consider long-duration bonds. Yue (2010) and Benjamin and Wright (2009) take renegotiations and settlements into consideration. All these papers examine equilibrium outcomes under centralized borrowing. Our paper instead studies outcomes under decentralized borrowing.

The remainder of the paper is organized as follows. Section 1.2 presents the model with decentralized borrowing. In section 1.3, we compare the quantitative implications of the models with decentralized and centralized borrowing. Section 1.4 investigates how different default penalties affect the quantitative results, in particular, the equilibrium debt level. We conclude in section 1.5.

[4] Interest rates depend on aggregate debt in the decentralized borrowing model. By contrast, one can interpret that interest rates depend on individual borrowing in the centralized borrowing model, which generates the same outcomes as a model with both decentralized borrowing and decentralized default.

1.2 Models

This section presents a dynamic stochastic general equilibrium model of decentralized borrowing and centralized default in which borrowing decisions are made by individual households and default decisions are made by a government. This setup is intended to capture an environment in which borrowing decisions are made by private agents and lending decisions of foreign lenders are guided by aggregate indicators rather than individual borrowers' ability to repay. By comparing with the centralized borrowing model this section highlights the pecuniary externality arising from decentralized borrowing.

1.2.1 Model with Decentralized Borrowing

The model economy consists of three types of agents: a continuum of identical households and a sovereign government in a small open economy, and foreign lenders. The households receive stochastic aggregate income shocks y, which follow a Markov process with the transition function $f(y', y)$. In order to smooth income shocks, the households trade non-contingent bonds b with risk-neutral foreign lenders. The benevolent government, maximizing its representative household's welfare, decides whether to enforce foreign debt contracts.[5] In each period, the country is either in the normal phase with access to international financial markets or in the penalty phase without access to financial markets.

The timing is as follows. At the beginning of each period, the income shock y is realized. If the country is in the normal phase, the government decides whether to

[5]Broner and Ventura (2010) analyze an environment with both domestic and international trade of contingent claims among private agents. They assume that the government, when deciding whether to enforce the claims, cannot discriminate between domestic and foreign creditors. We implicitly allow discrimination. Our framework is equivalent to the one in which both domestic and international borrowing and lending are allowed, and the government always enforces domestic contracts, but not necessarily international contracts. Given that households are identical, domestic borrowing and lending are never observed in equilibrium.

enforce the repayment of outstanding foreign debt B.[6] If the government enforces debt contracts, the households repay their debt b and decide on consumption c and next-period debt b'. If the government defaults, the households do not to repay their debt, and the economy goes into the penalty phase. The country in the penalty phase suffers from income loss and has probability θ of reverting to the normal phase each period.

Government

At the beginning of the normal period, the benevolent government observes current income shock y and aggregate foreign debt B. The government decides whether to enforce debt contracts to maximize the representative household's welfare. This welfare is given by $v^D(y)$ if the government chooses to default, and $v^R(B, y, \Gamma(B, y))$ if the government chooses to enforce the repayment with an anticipation that the economy will borrow $B' = \Gamma(B, y)$ this period. Thus, the government solves the following problem:

$$D(B, y) = \arg \max_{d \in \{0,1\}} \left\{ (1-d) \; v^R(B, y, \Gamma(B, y)) + d \, v^D(y) \right\}, \tag{1}$$

where $d = 1$ indicates default and $d = 0$ indicates repayment. If the repayment welfare v^R is greater than the default welfare v^D, then the government enforces the repayment of individual debt contracts. Otherwise, the government decides to declare default. Our assumption that national governments make default choices highlights default risk, driven by national governments, of private debt contracts. The governments can impose exchange or capital controls to prevent private agents from repaying their debt or assume repayment responsibilities to foreign creditors by nationalizing private foreign debt.

[6] A positive B denotes foreign assets, and a negative B denotes foreign debt.

Foreign Lenders

Foreign lenders are risk neutral. They operate in competitive international financial markets and have the opportunity cost of funds at the risk-free interest rate r. They thus have to break even for each debt contract. Since the government's default decisions are based on aggregate debt, the bond price schedule also depends on aggregate debt. For any aggregate borrowing level B', the lender expects to receive the repayment B' next period if and only if the government enforces repayment next period, that is $D(B', y') = 0$. Thus, the total expected repayment next period is $\int_{y'} B' \left(1 - D(B', y')\right) f(y', y) dy'$. The resource cost of this debt contract to the lender today is $q(B', y)B'$. The zero profit condition requires that the resource cost equals the present value of the expected repayment. This gives rise to the bond price schedule:

$$q(B', y) = \frac{\int_{y'} \left(1 - D(B', y')\right) f(y', y) dy'}{1 + r}. \tag{2}$$

If the government will enforce repayment under all future income shocks, the bond price is simply the inverse of the gross risk-free rate. However, if the government defaults for some future income shocks, the bond price is lower to compensate for the default risk.

The centralized default decision for decentralized borrowing implies that credit terms for private debt depends critically on aggregate debt. This implication is consistent with empirical evidence. Using firm-level observations from 30 countries for 1995–2004, Borensztein et al. (2007) show that sovereign ratings are a significant determinant of credit ratings assigned to corporations in emerging market economies. Similar findings are presented in Agca and Celasun (2009), Ferri and Liu (2003), Fernandez-Arias and Lombardo (1998), and Mendoza and Yue (2010). Major rating agencies' practices also support this implication. For example, Standard and Poor's

(2001) stresses that sovereign credit risk is always a key consideration in the assessment of the credit standing of banks and corporations in emerging markets. Their main argument is that governments in financial distress or default may force private sector to default by imposing exchange controls and other restrictive measures.

Individual Households

We now describe the individual household's problem. A measure one continuum of infinitely-lived identical households have flow utility $u(c)$ over consumption c, where $u(\cdot)$ is increasing and strictly concave. If the country is in the normal phase and the government decides to repay, then the households can trade one-period non-contingent bonds b'. The households take as given the aggregate borrowing level B' and the associated bond price $q(B', y)$. In addition, the households also take as given the default decision of the government $D(B', y')$.

Hence, a household with bond holding b and income shock y solves:

$$\begin{aligned} v^R(b, y, B') \quad = \quad & \max_{b'} \quad u(y + b - q(B', y)b') \\ & +\beta \int_{y'} \left[(1 - D(B', y')) \, v^R(b', y', B'') + D(B', y') \, v^D(y') \right] f(y', y) dy' \\ & s.t. \quad B'' = \Gamma(B', y'), \end{aligned} \tag{3}$$

where $0 < \beta < 1$ is the discount factor, and $B'' = \Gamma(B', y')$ is aggregate bonds that the economy will issue next period if the government continues to enforce repayment. Aggregate borrowing B' plays an important role in each household's decision; it pins down the cost of borrowing today, the government's default decision next period, and future aggregate borrowing B''.

If the government decides to default, the households do not repay their debt but lose access to international financial markets. In each period, the economy has probability θ of regaining access to international financial markets with zero debt obligations. During the exclusion periods, the households suffer from income loss; their income drops from y to y^{def}. The default welfare is given by

$$\begin{aligned} v^D(y) &= u(y^{def}) + \beta \int_{y'} \left[\theta v^R(0, y', B'') + (1-\theta) v^D(y')\right] f(y', y) dy', \\ & \quad s.t. \quad B'' = \Gamma(0, y'). \end{aligned} \tag{4}$$

Recursive Competitive Equilibrium

The recursive competitive equilibrium of this economy is a list of (i) individual value functions and policy functions: v^R, v^D, c, and b', (ii) a government default decision function $D(B, y)$, (iii) an actual law of motion for aggregate debt $B' = \Gamma(B, y)$, and (iv) a bond price schedule $q(B', y)$ such that

1. Given q, Γ and D, the value and policy functions solve the household's problem.
2. The household's policy function b' is consistent with Γ.
3. Given Γ, $D(B, y)$ solves the government's problem.
4. The bond price schedule $q(B', y)$ ensures foreign lenders' break-even in expected value.

1.2.2 Model with Centralized Borrowing

We compare our decentralized borrowing model with the standard Eaton and Gersovitz (1981) type model of centralized borrowing. All aspects of the model with centralized borrowing are identical to the model with decentralized borrowing, except one difference. In the centralized borrowing model, the government, instead of the households, makes the borrowing decision.[7] We briefly describe the centralized borrowing

[7] One can alternatively think of the government's problem as the representative household's problem.

model. The government's value function is

$$W(B,y) = \max_{d\in\{0,1\}} \left\{(1-d)W^R(B,y) + dW^D(y)\right\} \tag{5}$$

where $W^R(B,y)$ is the repayment welfare and $W^D(y)$ is the default welfare. Let $D_C(B,y)$ denote the government optimal default decision and $q_C(B',y)$ denote the bond price schedule. The repayment welfare is given by

$$W^R(B,y) = \max_{B'} \; u(y + B - q_C(B',y)B') + \beta \int_{y'} W(B',y')f(y',y)dy'. \tag{6}$$

Note that the government chooses aggregate debt next-period, B', and allocates it evenly across the households. The default welfare is defined as

$$W^D(y) = u(y^{def}) + \beta \int_{y'} \left[\theta W(0,y') + (1-\theta)W^D(y')\right] f(y',y)dy'. \tag{7}$$

The bond price schedule is again given by foreign lenders' break-even condition:

$$q_C(B',y) = \frac{\int_{y'} \left(1 - D_C(B',y')\right) f(y',y)dy'}{1+r}. \tag{8}$$

The recursive competitive equilibrium of this economy consists of a list of the government's value functions, $\{W, W^R, W^D\}$, policy functions $\{B',\ D_C\}$ and a bond price schedule $q_C(B',y)$ such that

1. Under q_C, the value and policy functions solve the government's problem.

2. The bond price $q_C(B',y)$ ensures foreign lenders' break-even in expected value.

1.2.3 Comparison of the Two Borrowing Environments

In order to facilitate exposition, we treat the value functions and the bond price functions as differentiable in this subsection.[8] The first order condition in the model with centralized borrowing is

$$u'(c)\left[q_C(B',y)+\frac{\partial q_C(B',y)}{\partial B'}B'\right]=\beta\int_{y'}\left(1-D_C(B',y')\right)u'(c')f(y',y)dy'. \tag{9}$$

The $\frac{\partial q_C(B',y)}{\partial B'}B'$ term represents the change in the bond price in response to one extra unit of the bond. This term is not present in the corresponding first order condition in the model with decentralized borrowing:

$$u'(c)q(B',y)=\beta\int_{y'}\left(1-D(B',y')\right)u'(c')f(y',y)dy', \tag{10}$$

since households take the bond price as given.

For expository purposes, let us compare the debt levels assuming that the bond price schedules and the default sets are the same in the two models, that is $q_C = q$ and $D_C = D$. Denote the optimal bond holdings in the model with centralized borrowing and in the model with decentralized borrowing by B'_C and B'_D, respectively.[9] For sufficiently low levels of debt, the government enforces repayments under all future shocks and thus the economy faces the risk-free interest rate. We denote the maximum amount of such debt by $\overline{B}'$ and refer to it as the safe debt limit. Then it must be the case that $\frac{\partial q(B',y)}{\partial B'}=0$ for any $B'>\overline{B}'$. This implies that $B'_C = B'_D$ if the optimal debt is below the safe debt limit in both models.

Now consider the effect of raising debt by one unit when $B'<\overline{B}'$. The marginal cost is the expected loss in future utility conditional on not defaulting next period,

[8]The solution method employed in the quantitative analysis section does not depend on the differentiability of the value functions and the bond price schedule.

[9]With decentralized borrowing, individual households choose b'_D instead of B'_D. In equilibrium, however, individual and aggregate debt coincide. Thus, we compare aggregate debt in the two models.

which is the right hand side of equation (9) and (10). The marginal benefit is the current utility gain from the resource raised by one extra unit of debt, which is the left hand side of these two equations. We plot the marginal cost and benefit for each model in Figure 2. The marginal costs are identical across the two models and rise with the debt level.[10] The marginal benefits in both models decline with the debt level. Moreover, the marginal benefit is higher under decentralized borrowing since $\frac{\partial q(B',y)}{\partial B'} > 0$ and $\frac{B' \partial q(B',y)}{\partial B'} < 0$. At the optimal debt level, the marginal benefit equals the marginal cost. This implies that $B'_C > B'_D$, and so the households would like to borrow more under decentralized borrowing.

Figure 2: Marginal Benefits and Marginal Costs of Debt

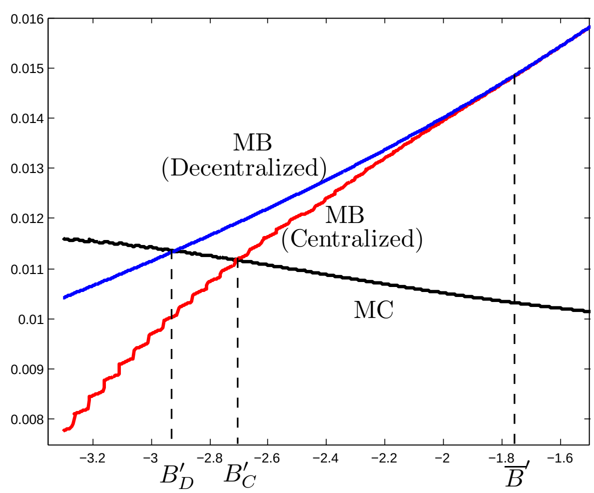

Notes: The marginal cost and benefit of an additional unit of debt in the centralized borrowing model are plotted for a country with an endowment shock at the 60th percentile and a debt level about 34% of income using the model solution. The marginal benefit in the decentralized borrowing model is constructed using the bond price schedule in the centralized borrowing model.

When making borrowing decisions, the government internalizes the adverse effect of additional borrowing on the bond price, but individual households, acting as price takers, do not.[11] Thus, decentralized borrowing generates a pecuniary externality

[10]In principle, marginal costs might decrease with debt if default probabilities rise rapidly with debt.

[11]The externality resembles the one studied in Gertler and Kiyotaki (2010). In their paper, individual banks do not internalize the effects of their own liability structure on the aggregate liability structure. This externality induces relatively more debt financing in individual banks' liability structure, and leads to an over-levered aggregate balance sheet.

where one individual's actions affect another individual's welfare through prices.[12] Pecuniary externalities by themselves are not a source of inefficiency since they work within the market mechanism through prices. However, they do cause efficiency losses and lower welfare if there are other market imperfections such as incomplete markets and limited enforcement in the model.[13]

The above discussions assume that the bond price and default schedules are the same in the two models. These assumptions automatically hold if default never occurs in equilibrium and the bond price schedule is an exogenous function of aggregate debt. In this case, decentralized borrowing unambiguously leads to overborrowing in equilibrium.[14] However, in our model both the bond price and default schedules are endogenous. Given the overborrowing incentives of the households, borrowing costs are higher and welfare, especially the repayment welfare, is lower under decentralized borrowing. Consequently, the government has a higher incentive to default, and the bond price schedule is less favorable under decentralized borrowing, i.e., the default set changes and the bond price schedule shifts. This bond price schedule effect reduces borrowing. Hence, whether decentralized borrowing leads to equilibrium overborrowing depends on which effect dominates: the overborrowing incentive or the bond price schedule effect. We analyze quantitatively the impacts of decentralized borrowing on the equilibrium debt level in the next section.

1.3 Quantitative Analysis

This section investigates the quantitative implications of the decentralized borrowing model. In order to highlight the impacts of decentralized borrowing, we first

[12]Levchenko (2005) highlights another source of externality of private borrowing. When there are heterogeneous agents and heterogeneous access to international financial markets, financial integration might break domestic risk sharing and hurt those without access to international financial markets.

[13]For more discussions on efficiency losses from pecuniary externalities, see Loong and Zeckhauser (1982) and Greenwald and Stiglitz (1986).

[14]This is one of the examples in Uribe (2006) and referred to as the debt-elastic country premium case.

compare the equilibrium dynamics of the decentralized and centralized borrowing environments. We then evaluate the ability of the decentralized borrowing model to account for observed statistical moments of the business cycle in Argentina.

1.3.1 Calibration and Computation

We calibrate the model at the quarterly frequency. The utility has standard CRRA form: $u(c) = \frac{c^{1-s}-1}{1-s}$, where the coefficient of relative risk aversion s is 2. The risk-free interest rate is set to 1.7%, corresponding to the average quarterly interest rate of a five-year U.S. treasury bond for the period 1983–2001. The income shock y_t follows an AR(1) process: $\ln(y_t) = \rho \ln(y_{t-1}) + \varepsilon_t$ with $|\rho| < 1$ and $\varepsilon_t \sim N(0, \sigma_\varepsilon^2)$. We use the time series of Argentina's GDP to calibrate the shock process and estimate ρ to be 0.945 and σ_ε to be 0.025.

The default penalty plays a crucial role in sovereign debt models. In the benchmark calibration, we assume that the default penalty is disproportionately large for large income shocks, following Arellano (2008). Specifically, y^{def} has the following form:

$$y^{def} = \begin{cases} (1-\lambda)\bar{y} & \text{if } y > (1-\lambda)\bar{y} \\ y & \text{if } y \leq (1-\lambda)\bar{y} \end{cases}, \tag{11}$$

where $\bar{y}$ denotes the unconditional mean of income shocks, and λ characterizes the income loss after default. A larger λ makes the default penalty more severe both by lowering the threshold income shock that is subject to income loss and by raising the magnitude of income loss. We refer to this specification as the *asymmetric default penalty*. An alternative specification is the *symmetric default penalty* where income loss is a constant fraction of the income shock.

The motivation for the asymmetric default penalty is that sovereign default is often accompanied by a drop in private credit, and so the economy would have to forgo disproportionately larger income under good shocks.[15] In addition, with a symmetric default penalty, sovereign debt models rarely generate equilibrium default and fail to match the default rate observed in the data. The asymmetric default penalty makes default attractive when the country experiences bad shocks, and thus helps raise the default probability. Empirical support on either form of the default output cost is rather weak despite their common use. Given the quantitative importance of the default output cost, we conduct a sensitivity analysis on the default penalties in the next section.

The default penalty parameter λ, the discount factor β, and the re-entry probability θ are chosen such that the model with decentralized borrowing produces the 3% default probability, 14% income drop upon default, and 1.75% standard deviation of the trade balance observed in the Argentina data. The default penalty parameter λ is estimated to be 0.1 and the discount factor β is 0.97. The re-entry probability θ is estimated to be 0.1, which corresponds to 10 quarters of exclusion from international financial markets after default. This is in line with the historical evidence presented in Gelos et al. (2010).[16] See the lower panel of Table 1 for the summary of these parameter values.

With the functional forms and parameters described above, we solve the models numerically using the discrete state-space technique. The decentralized borrowing model is more difficult to compute than the centralized borrowing model. The state space has three dimensions $(b, y; B')$ in the decentralized borrowing model, while it has only two dimensions (B, y) in the centralized borrowing. Moreover, in the decentralized borrowing model we need to find the aggregate borrowing function $\Gamma(B, y)$ to be consistent with individual borrowings. Such aggregate borrowing functions are not unique in general. We discuss the detailed solution algorithm and the selection

[15] Mendoza and Yue (2010) present a model that generates endogenously this form of income loss.

[16] Gelos et al. (2010) find, for all defaulting episodes during the period of 1980–2000, that the median exclusion length is 3 years after default.

method for the aggregate borrowing function in Appendix A.

After solving the model, we simulate the model for 500,000 periods and find the latest 1,000 default episodes. We extract 74 consecutive observations of the normal period before each default event and examine the mean statistics over these samples. The 74 observations prior to a default episode correspond to the number of quarters between the latest two default events in Argentina.[17] In the next subsection, we compare the implications of the decentralized borrowing model with those of the centralized borrowing model.

1.3.2 Decentralized versus Centralized Borrowing

Table 1 presents statistics for the Argentina data and for the decentralized and centralized borrowing models. The first column shows business cycle statistics for Argentina from 1983 to 2001. The annual default probability of 3% is based on three default episodes in approximately one hundred years. The average debt over GDP ratio of 43.36% is calculated for the period from 1983 to 2001 using *Global Development Finance*. The debt statistics include total external debt of the private and public sectors. The second column of Table 1 presents the statistics in the model with decentralized borrowing. To highlight the role of decentralized borrowing, we present these statistics in the model with centralized borrowing under the same set of parameter values in the third column.

There are three striking differences between the decentralized and centralized borrowing models. First, the mean spread under decentralized borrowing is higher by a factor of more than thirty compared to the one under centralized borrowing: 11.25% versus 0.37%. Second, the model with decentralized borrowing exhibits a much higher default probability, 3.03%, far exceeding 0.11% in the model with centralized borrowing. Third, the decentralized borrowing model generates a higher debt to income

[17] In the model with centralized borrowing, default is so rare that only 137 samples satisfy our criteria. We thus compute the model statistics based on these 137 samples.

Table 1: Comparison of Decentralized and Centralized Borrowing

Statistics	Data	Model		
		Decentralized	Centralized 1	Centralized 2
mean(spread)	10.31	11.25	0.37	7.30
std(spread)	5.60	26.09	0.76	7.63
mean(B/y)	−43.36	−22.48	−21.22	−7.23
std(y)	8.12	5.52	5.63	5.71
std(C)	9.47	6.57	6.35	6.38
std(TB/y)	1.75	1.75	1.49	1.75
corr(c, y)	0.98	0.97	0.97	0.96
corr(TB, y)	−0.59	−0.47	−0.36	−0.24
corr(spread, y)	−0.89	−0.55	−0.50	−0.67
corr(spread, c)	−0.91	−0.70	−0.64	−0.73
corr(spread, TB/y)	0.70	0.85	0.83	0.47
				0.00
prob(default)	3.00	3.03	0.11	3.07
drop in y upon default	−14.21	−13.28	−15.41	−13.02
drop in c upon default	−16.01	−12.77	−15.02	−12.84
welfare		1.000	1.008	1.013
welfare with mean y, zero b		1.000	1.007	1.011
discount factor β		0.97	0.97	0.93
output loss λ		0.10	0.10	0.10
re-entry probability θ		0.10	0.10	0.65

Note: The first column shows statistics for Argentina from 1983 to 2001. The income and consumption drops in default are based on the 2001 Argentine default episode. The interest rate spread is computed as the difference of the EMBI yield and the yield of a 5 year U.S. bond. The second column presents the statistics in the model with decentralized borrowing. The third column presents the statistics in the model with centralized borrowing under the same set of parameter values as in the decentralized borrowing model. The last column presents the statistics in the centralized borrowing model recalibrated to best match the data. All statistics except correlations and welfare are in percentage terms. The welfare results are calculated in terms of permanent consumption, and then normalized by the welfare level in the decentralized borrowing model for ease of comparison.

ratio than the centralized borrowing model does. The mean debt to income ratio is 22.48% in the decentralized borrowing model, while it is 21.22% in the centralized borrowing model.

To understand these differences, we examine borrowing decisions in the two models. Figure 3 plots the desired borrowing conditional on not defaulting over the current bond holdings.[18] Desired borrowing is similar across the two models for low levels of debt. As the debt level increases, desired borrowing increases faster under decentralized borrowing. Under centralized borrowing, the government recognizes that the interest rate increases as an additional unit of debt is taken. Under decentralized borrowing, however, households do not take into account the interest rate effect of their borrowing and would like to overborrow. This negative externality becomes especially severe when current debt is large and the interest rate rises sharply with an additional unit of debt.

Figure 3: Comparison of Desired Borrowing

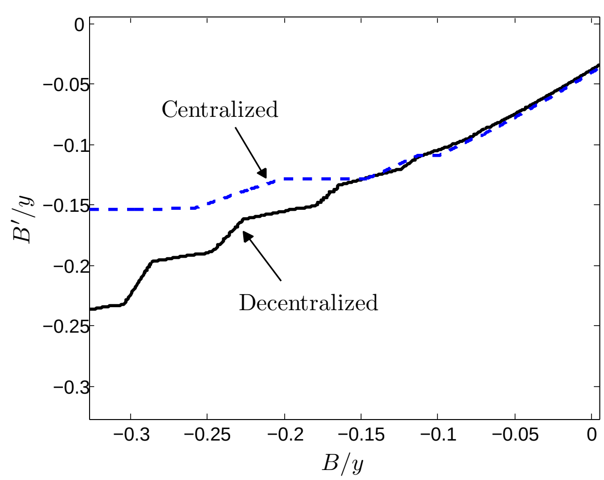

Borrowing and default are two instruments with which households affect their consumption path. Under centralized borrowing, the government, or equivalently the representative household, owns both instruments. Under decentralized borrowing, the

[18]All figures in this subsection are based on the income shock, which is 10% below the mean. We observe the same qualitative results for the other income shocks. Both the current and next-period bond holdings are normalized by the mean income shock.

households have only the first instrument and tend to take more debt since they fail to internalize the negative externality of their borrowing. Thus, welfare, especially the repayment welfare, is lower under decentralized borrowing, as shown in Figure 4. Consequently, the government finds default attractive for a wider range of debt levels under decentralized borrowing.

Figure 4: Comparison of Value Functions

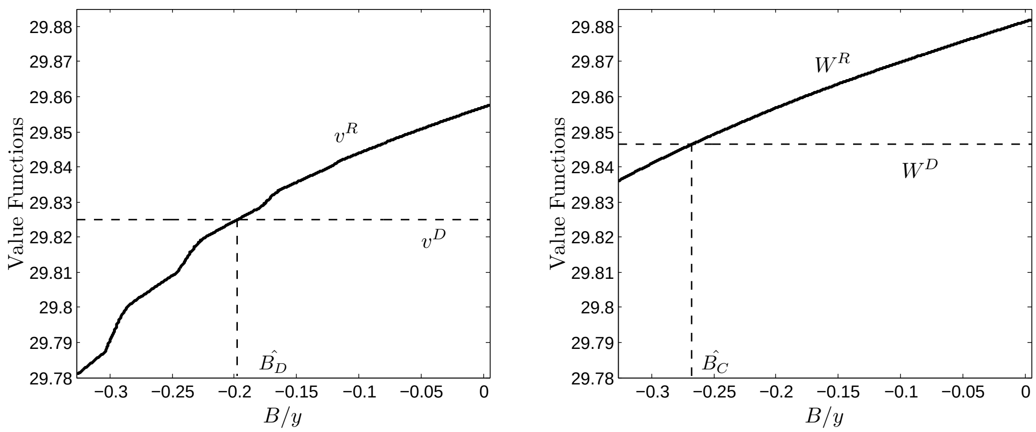

(a) Decentralized Borrowing

(b) Centralized Borrowing

The failure of individual households to internalize the effect of their borrowing on the government's default choices lowers the bond price schedule under decentralized borrowing. As shown in the left panel of Figure 5, the prices are discounted by more for any level of bonds under decentralized borrowing. This less favorable bond price schedule generates tighter debt limits and constrains borrowing by more. The right panel displays the total resources that debt generates, $q(B', y)B'$. Any debt B' less than the safe debt limit $\overline{B}'$ generates resources $B'/(1 + r)$. Once it exceeds the safe debt limit, the debt becomes risky. Let us refer to the debt level which maximizes the resource obtained from foreign lenders, $q(B', y)B'$, as the risky debt limit, denoted by $\underline{B}'$. The optimal level of debt would never exceed the risky debt limit because the borrower can obtain the same amount of resources with a smaller next-period repayment. As shown in the figure, both the safe and risky debt limits are tighter under decentralized borrowing.

Figure 5: Comparison of Bond Price Schedules

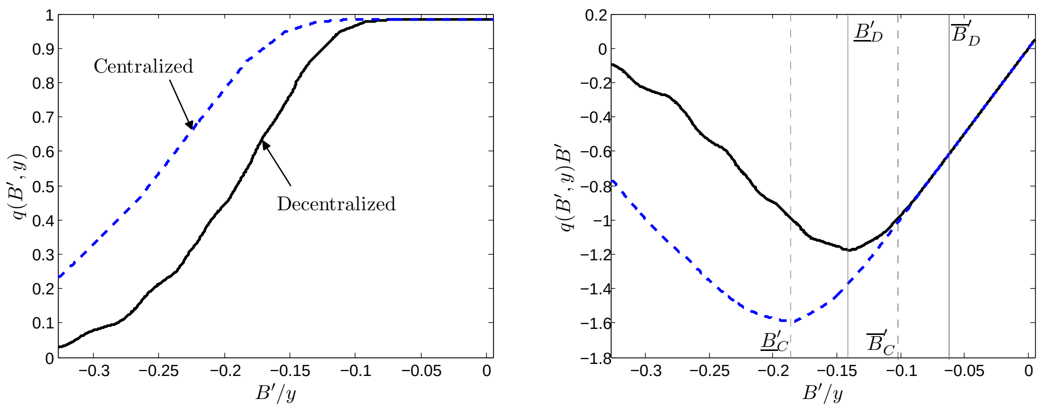

(a) Bond Price Schedule (b) Total Resource Borrowed

The incentive to overborrow and the lower bond price schedule have opposite effects on the equilibrium level of debt. Whether decentralized borrowing leads to larger equilibrium debt depends on which force dominates. Under the benchmark calibration, desired borrowing is higher even when the bond price schedule is less favorable under decentralized borrowing. This leads to higher equilibrium debt under decentralized borrowing. Figure 6 shows the limiting distribution of bond holdings as shares of mean income for the two models. The distribution is more concentrated on high debt levels in the decentralized borrowing model. This implies that even with higher costs of borrowing, the incentive to overborrow is strong enough to induce the households to issue more debt. As a result, the interest rate spread is substantially higher under decentralized borrowing. Also, the interest rate spread is more countercyclical under decentralized borrowing. As shown in Table 1, corr(spread, y) and corr(spread, c) are -0.55 and -0.70, respectively, under decentralized borrowing, and -0.50 and -0.64, respectively, under centralized borrowing.

Table 1 reports two sets of welfare statistics for each simulated model, first based on the limiting distribution and second holding debt and income constant at zero debt and at the mean income level across economies. The welfare is 1% lower under decentralized borrowing than under centralized borrowing. The 1% welfare difference

Figure 6: Comparison of Bond Distributions

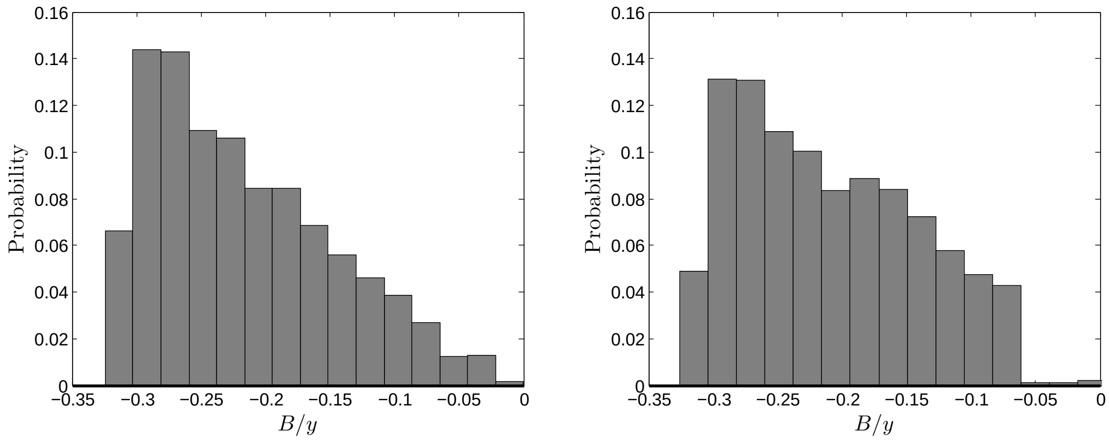

(a) Decentralized Borrowing

(b) Centralized Borrowing

is economically significant, considering that the welfare cost of business cycles estimated by Lucas (1987) is only about one-tenth of a percent of consumption. This welfare implication holds even when we compare the two models for any given level of the income shock and the debt-to-income ratio. The magnitudes of welfare differences vary from 0.6% to 1.2%.

In summary, the decentralized borrowing model generates a larger default rate, a higher mean spread, lower welfare and larger equilibrium debt than the centralized borrowing model in the benchmark calibration.[19] All these findings, except the one on equilibrium debt, are robust to different default penalty parameters and specifications. We will focus on the effects of decentralized borrowing on equilibrium debt levels in section 1.4.

[19]Models with political economy can also generate interest rate spreads higher than the standard sovereign debt model (see Cuadra and Sapriza (2008)). Political instabilities lead to short-sighted governments, who do not fully internalize the *next-period marginal cost* of their borrowing and have large incentives to default. Thus, everything else equal, interest rate spreads are higher than the one without political instabilities. By contrast, our model generates higher interest rate spreads because households do not fully internalize the *current-period marginal benefit* of their borrowing. They ignore the impact of an additional unit of debt on economy-wide borrowing costs and tend to overborrow, which increases default incentives of the government and interest rate spreads charged by international creditors.

1.3.3 Quantitative Predictions of the Models

In this subsection, we compare the quantitative predictions of the two models with the Argentine data. Both models are calibrated to match the relevant moments in the data. In particular, the parameters β, λ, and θ are calibrated to best match the default rate, income loss, and trade-balance volatility in the data. The fourth column of Table 1 shows the statistics of the recalibrated centralized borrowing model. To generate the data moments, the centralized borrowing model needs a low discount factor of 0.93, and a large reentry probability of 0.65. That is, the average exclusion period after default is only 1.5 quarters.

The most striking difference across these two models is the equilibrium debt level. The debt to income ratio is 7.23% in the model with centralized borrowing. In contrast, it is about 22.48% in the model with decentralized borrowing—much closer to the data. We want to highlight that these quantitative results depend critically on the form and the parameter values of the default penalties, which we will analyze in details in the next section. Given the importance of the default penalties, the literature will benefit from more empirical research on the default penalties.

The model with decentralized borrowing also shows better performance in terms of replicating countercyclical trade balances. The correlation between the trade balance and income is -0.59 in the data; it is only -0.24 in the model with centralized borrowing, but -0.47 in the model with decentralized borrowing. In addition, decentralized borrowing generates a mean spread close to the data of 10.31%. The mean spread is 11.25% in the decentralized borrowing model and only 7.30% in the centralized borrowing model. On the other hand, the decentralized borrowing model overestimates the volatility of the interest spread.

Although both models match well the output drop after default in the Argentine 2001 default episode, they predict a much stronger negative relationship between output and default than is found in the historical record. As documented by Tomz and Wright (2007), the output declines from trend by only 1.6% in the first year of default in 169 default episodes over the period 1820–2004. Moreover, they document that only 62% of all the default episodes occur when the output is below the trend. However, the output declines from the trend in the first year of default by 12% under decentralized borrowing and by 11% under centralized borrowing. In both models, almost all default episodes occur when the output is below the trend.[20]

One caveat of the model simulation is worth mentioning. We compute model statistics based on 74 simulation periods before default to mimic the 74 quarters between the two default events of Argentina. Ideally, we should compute statistics based on the simulations where two consecutive defaults are exactly 74 periods apart. However, it is rare to find such cases in the simulation. Instead, we compute the model statistics based on the simulation episodes in which default ends and starts 70–78 periods apart, and compare these statistics with the benchmark results in Table 4 in the appendix.[21] Although some statistics are moderately different from the benchmark, the key patterns between the centralized and decentralized borrowing models are robust to this restricted sample of the model simulations.

1.4 Overborrowing or Underborrowing

This section examines the quantitative effect of decentralized borrowing on equilibrium levels of debt for different default penalties. We find that decentralized borrowing generates larger levels of equilibrium debt than centralized borrowing only when the default penalties are asymmetric and lenient. The decentralized borrowing model generates low debt levels under symmetric default penalties or under asymmetric harsh default penalties.

[20] A lower direct output cost parameter would reduce the discrepancy between the model simulation and the data in these dimensions in both models. However, the models would have difficulties matching the observed frequency of default and the observed debt to output ratio.

[21] There are about 13 such episodes in a simulation of 500,000 periods.

1.4.1 Alternative Default Penalty Parameters

Consider the two models with the benchmark parameter values. In the first set of experiments, we vary the default income loss parameter λ from 1% to 20% while fixing all the other parameters. In the second set of experiments, we vary the re-entry probability from 1% to 40% while fixing all the other parameters. We plot the equilibrium debt to income ratios of the two models for these two sets of experiments in Figure 7. First of all, equilibrium debt in both models increases with the default income loss λ and decreases with the re-entry probability θ. This is intuitive because larger values of λ or lower values of θ are associated with more severe default penalties, and this in turn implies less frequent default and more lenient bond price schedules.

Figure 7: Equilibrium Debt: Varying Default Penalties

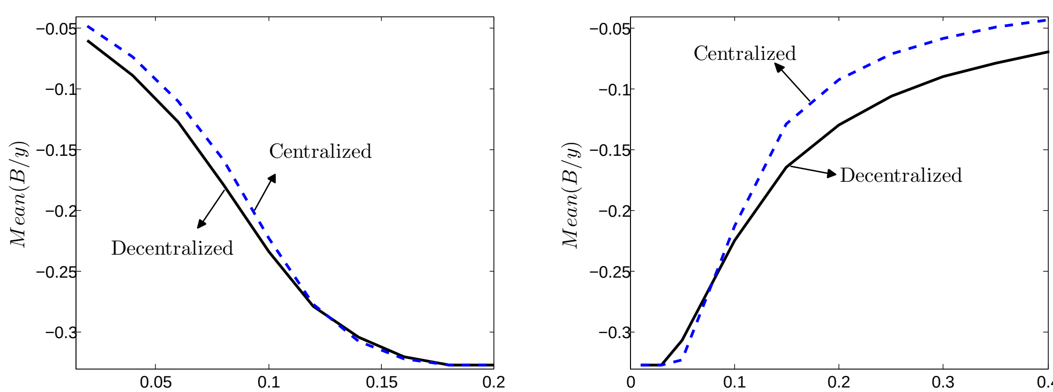

(a) Default Income Loss λ (b) Re-entry Probability θ

Second, we find that decentralized borrowing generates overborrowing for low values of λ, but underborrowing for high values of λ, as shown in the left panel of Figure 7. The differences in equilibrium debt appear to be small in the figure, but the magnitudes of overborrowing or underborrowing are not trivial. For example, decentralized borrowing generates overborrowing by 24.3% when λ is 0.02 and underborrowing by 1.2% when λ is 0.14. Note that for λ higher than 0.18, no default happens and thus the equilibrium debt levels are identical in both models.

As we discussed earlier, decentralized borrowing has two counteracting effects on equilibrium debt: the overborrowing and bond price schedule effects. The overborrowing incentive arises from the failure of individual households to internalize the effect of their additional borrowing on the bond price. The more sensitive credit costs are to aggregate borrowing, the stronger is the overborrowing effect, which can be seen from comparing the first order conditions in equation (9) and (10). On the other hand, the bond price schedule effect occurs because the overborrowing incentive lowers the repayment welfare and drives up the default likelihood of the government. Thus, the greater the decrease in the repayment welfare is, the stronger is the bond price schedule effect.

When λ is low, the bond price schedules in both models are very sensitive to aggregate risky borrowing because the government's incentive to default rises rapidly with additional units of debt. In addition, the difference between the two bond price schedules is small, as shown in the left panel of Figure 8. This is because decentralized borrowing, though it produces more defaults, lowers the repayment welfare only by a little because default is not that costly. Thus, the overborrowing effect is strong while the bond price schedule effect is weak, which leads to equilibrium overborrowing. As λ increases, the bond price schedules in both models become flatter and the difference between them becomes larger, as shown in the right panel of Figure 8. With increasingly severe default penalties, the government's incentive to default does not rise quickly as aggregate debt rises. But decentralized borrowing reduces the repayment welfare substantially, which drives up the default incentives and lowers the bond price schedule greatly. Thus, the overborrowing effect is weak while the bond price schedule effect is strong, which leads to equilibrium underborrowing.

Finally, we compare equilibrium debt of the two models for different re-entry probabilities θ. Decentralized borrowing generates overborrowing when θ is high, but underborrowing when θ is low. In particular, equilibrium debt under decentralized borrowing is 53.1% more when θ is 0.3, but 5.3% less when θ is 0.05. The equilibrium

Figure 8: Bond Prices for Different Income Losses

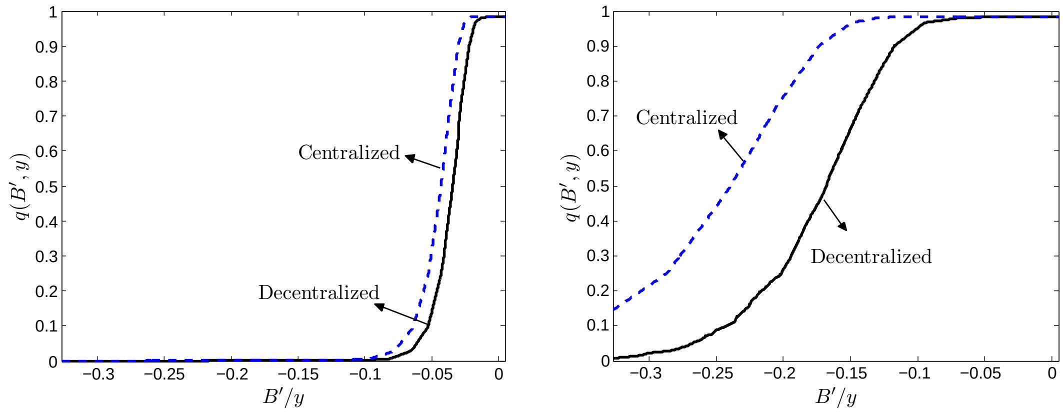

(a) Low Income Loss ($\lambda = 0.04$) (b) High Income Loss ($\lambda = 0.14$)

debt levels are identical across the two models for very low values of θ, which implies no default in equilibrium. The intuition for these results is similar to that for different λ. When θ is high, the bond price schedules are steep and similar in both models, which leads to equilibrium overborrowing. When θ is low, the bond price schedules become flatter and the difference between them becomes larger, implying equilibrium underborrowing.

When the default rates are zero in both models, the two models generate identical business cycle statistics including the mean spread, default rate and welfare. In all the other cases, the model with decentralized borrowing generates larger default rates, higher mean spreads, and lower welfare than the model with centralized borrowing. Even when decentralized borrowing generates equilibrium underborrowing, the substantial difference in the bond price schedule, as shown in the right panel of figure 8, is enough to make the spreads higher in equilibrium. We report the detailed statistics for these experiments in Appendix B.

1.4.2 Alternative Default Penalty Specification

We now investigate the symmetric default penalty of the following form:

$$y^{def} = (1 - \eta)y, \text{ for all } y, \tag{12}$$

where η captures the constant fraction of income loss after default. Figure 9 shows equilibrium debt levels for different values of η and θ. Surprisingly, under the symmetric default penalty, decentralized borrowing consistently generates underborrowing.

Figure 9: Equilibrium Debt under Symmetric Default Penalty

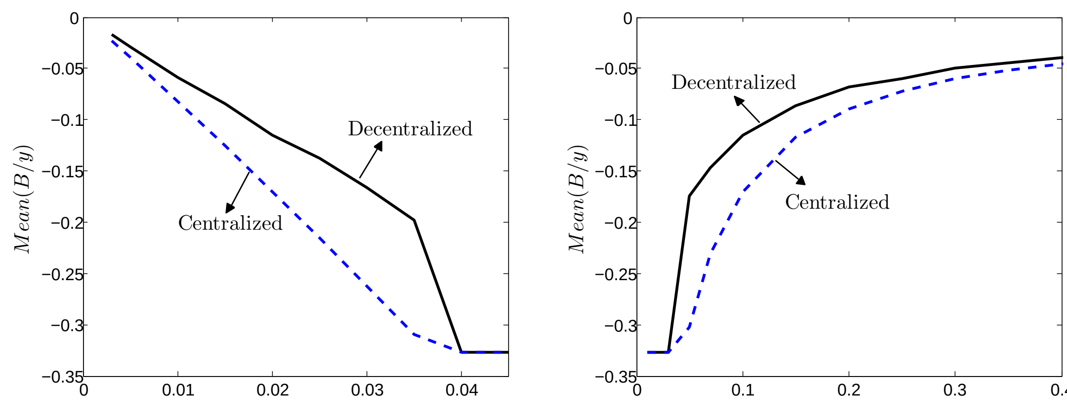

(a) Default Income Loss η (b) Re-entry Probability θ

To understand this result, we plot the bond price schedules for both models under the symmetric default penalty in the left panel of Figure 10. In the right panel, we plot the resources obtained from foreign lenders, $q(B', y)B'$, as a function of next-period debt. The bond price schedules are extremely steep and the risky-debt region $[\underline{B}', \overline{B}']$ is tiny in both models. For the overborrowing effect to operate, the risky debt region needs to be large to accommodate the overborrowing incentives. Given the tiny risky debt region, equilibrium debt is mainly constrained by the safe debt limits in both models. Under decentralized borrowing, the overborrowing incentives tighten the safe debt limit greatly, and this translates into underborrowing in equilibrium.

As in the case with the asymmetric default penalty, the model with decentralized

Figure 10: Bond Prices Under Symmetric Default Penalty

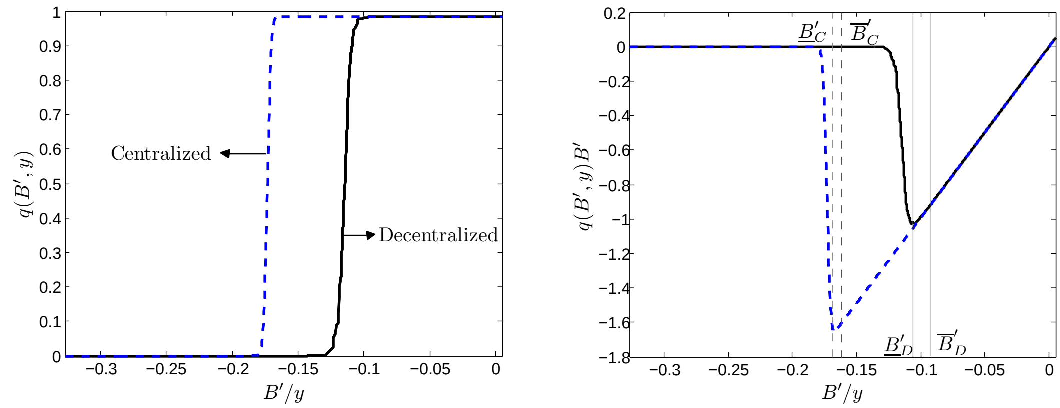

(a) Bond Price Schedule (b) Total Resource Borrowed

borrowing generates a higher default probability, a higher mean spread, and lower welfare for all the cases with positive default rates. We report the detailed statistics for different values of parameters in the symmetric default penalty case in Appendix B.

1.5 Conclusion

Private debt inflows to developing countries have risen substantially in the past two decades. Given the central role of developing countries' governments in foreign debt repayments, private foreign debt is often priced with macroeconomic indicators instead of individual borrowers' ability to repay. In such an environment, a pecuniary externality arises from decentralized borrowing because private agents do not internalize the negative impact of their borrowing on aggregate credit costs. It has been widely argued that this pecuniary externality caused excessive borrowing and frequent debt crisis in developing countries. This paper evaluated quantitatively the pecuniary effect of decentralized borrowing in a model where individual households make borrowing decisions and a government makes default decisions to maximize the welfare of the representative household.

Despite the overborrowing incentives, the model with decentralized borrowing generates a lower level of equilibrium debt than the model with centralized borrowing for a wide range of paramter values. This is because households also face a less favorable bond price schedule under decentralized borrowing, which tends to reduce the optimal level of debt. When the income loss is proportional to the income shock, decentralized borrowing always generates a lower equilibrium debt level regardless of the parameter values for default penalties. When the income loss after default is disproportionately large under good income shocks, decentralized borrowing generates underborrowing for severe default penalties, but overborrowing for lenient default penalties. On the other hand, decentralized borrowing unambiguously drives up the economy-wide credit costs, raises the likelihood of sovereign default, and lowers welfare.

Given our analysis on decentralized borrowing, regulations on private international capital flows may improve welfare. The most obvious policy would be imposing capital controls to prohibit private borrowing. This would require that the government be able to efficiently allocate funds among private agents. Alternatively, the government can impose, on international private borrowing, either taxes if default is not that costly or subsidies if default is costly.[22] Future research on the optimal tax or subsidy on international private borrowing will be useful since in practice it is hard to implement capital controls.

[22]For discussions of optimal policy under complete markets, see Jeske (2006), Kehoe and Perri (2004) and Wright (2006).

References

Agca, Senay and Oya Celasun, "How Does Public External Debt Affect Corporate Borrowing Costs in Emerging Markets?," *IMF Working Paper, 09/266*, 2009.

Aguiar, Mark and Gita Gopinath, "Defaultable Debt, Interest Rates, and the Current Account," *Journal of International Economics*, June 2006, *69* (1), 64–83.

Arellano, Cristina, "Default Risk and Income Fluctuations in Emerging Economies," *American Economic Review*, June 2008, *98* (3), 690–712.

_ **and Ananth Ramanarayan**, "Default and the Maturity Structure in Sovereign Bonds," *Working Paper*, 2010.

Bai, Yan and Jing Zhang, "Financial Integration and International Risk Sharing," *Working Paper*, 2009.

_ **and** _ , "Solving the Feldstein-Horioka Puzzle with Financial Frictions," *Econometrica*, 2010, *78* (2), 603–632.

Benjamin, David and Mark L.J. Wright, "Recovery Before Redemption: A Model of Delays in Sovereign Debt Renegotiations," *Working Paper*, 2009.

Bi, Ran, "Debt Dilution and the Maturity Structure of Sovereign Bonds," *Working Paper, University of Maryland*, 2006.

Bizer, David S. and Peter M. DeMarzo, "Sequential Banking," *Journal of Political Economy*, 1992, *100* (1), 41–61.

Borensztein, Eduardo, Kevin Cowan, and Patricio Valenzuela, "Sovereign Ceilings "Lite"? The Impact of Sovereign Ratings on Corporate Ratings in Emerging Market Economies," *IMF Working Paper, 07/75*, 2007.

Broner, Fernando and Jaume Ventura, "Globalization and Risk Sharing," *Review of Economic Studies (forthcoming)*, 2010.

Chatterjee, Satyajit and Burcu Eyigungor, "Maturity, Indebtedness, and Default Risk," *Working Paper*, 2010.

Cuadra, Gabriel and Horacio Sapriza, "Sovereign Default, Interest Rates and Political Uncertainty in Emerging Markets," *Journal of International Economics*, 2008, *76*, 78–88.

Eaton, Jonathan and Mark Gersovitz, "Debt with Potential Repudiation: Theoretical and Empirical Analysis," *Review of Economic Studies*, April 1981, *48* (2), 289–309.

Fernandez-Arias, Eduardo and Davide Lombardo, "Private External Overborrowing in Undistorted Economies: Market Failure and Optimal Policy," *Inter-American Development Bank Working Paper Series 369*, 1998.

Ferri, Giovanni and Li-Gang Liu, "How Do Global Credit-Rating Agencies Rate Firms from Developing Countries?," *Asian Economic Papers*, 2003, *2* (3), 30–56.

Gelos, R. Gaston, Ratna Sahay, and Guido Sandleris, "Sovereign Borrowing by Developing Countries: What Determines Market Access?," *Journal of International Economics (forthcoming)*, 2010.

Gertler, Mark and Nobuhiro Kiyotaki, "Financial Intermediation and Credit Policy in Business Cycle Analysis," *Working Paper*, 2010.

Greenwald, Bruce C. and Joseph E. Stiglitz, "Externalities in Economies with Imperfect Information and Incomplete Markets," *Quarterly Journal of Economics*, May 1986, *101* (2), 229–264.

Hatchondo, Juan Carlos and Leonardo Martinez, "Long-Duration Bonds and Sovereign Defaults," *Journal of International Economics*, 2009, *79* (1), 117–125.

_ , _ , and Horacio Sapriza, "Heterogeneous Borrowers in Quantitative Models of Sovereign Default," *International Economic Review*, 2009, *50* (4), 1129–1151.

_ , _ , and _ , "Quantitative Properties of Sovereign Default Models: Solution Methods Matter," *Review of Economic Dynamics*, 2010, *13* (4), 919–933.

Jeske, Karsten, "Private International Debt with Risk of Repudiation," *Journal of Political Economy*, June 2006, *114* (3), 576–593.

Kehoe, Patrick J. and Fabrizio Perri, "Competitive Equilibrium with limited enforcement," *Journal of Economic Theory*, 2004, *119* (1), 184–206.

Levchenko, Andrei, "Financial Liberalization and Consumption Volatility in Developing Countries," *IMF Staff Papers*, 2005, *52:2*, 237–259.

Loong, Lee Hsien and Richard Zeckhauser, "Pecuniary Externalities Do Matter When Contingent Claims Markets Are Incomplete," *Quarterly Journal of Economics*, Feb. 1982, *97* (1), 171–179.

Lorenzoni, Guido, "Inefficient Credit Booms," *Review of Economic Studies*, 2008, *75*, 809–833.

Lucas, Robert E. Jr., *Models of Business Cycles*, Oxford: Basil Blackwell, 1987.

Mendoza, Enrique G. and Vivian Z. Yue, "A Solution to the Disconnect between Country Risk and Business Cycles Theories," *Working Paper*, 2010.

Standard and Poor's, "Rating Methodology: Evaluating the Issuer, Corporate Ratings Criteria," September 2001.

Tauchen, George and Robert Hussey, "Quadrature-Based Methods for Obtaining Approximate Solutions to Nonlinear Asset Pricing Models," *Econometrica*, 1991, *59* (2), 371–396.

Tomz, Michael and Mark L.J. Wright, "Do Countries Default in 'Bad Times'?," *Journal of the European Economic Association*, April-May 2007, *5* (2-3), 352–360.

Uribe, Martin, "On Overborrowing," *American Economic Review, Papers and Proceedings*, May 2006, *96* (2), 417–421.

Wright, Mark L.J., "Private Capital Flows, Capital Controls and Default Risk," *Journal of International Economics*, June 2006, *69* (1), 120–149.

Yue, Vivian Z., "Sovereign Default and Debt Renegotiation," *Journal of International Economics*, 2010, *80* (2), 176–187.

Appendix A – Computational Algorithm

This appendix describes the computation algorithm for the decentralized borrowing model in details. We first discretize the state space (b, y, B'). We next make initial guesses for the bond price schedule and the government default decision. Specifically, we assume that $q^0(B', y) = \frac{1}{1+r}$ for all (B', y) and $d^0(B, y) = 0$ for all (B, y). Given these guesses, we solve the individual household's optimal debt level $b'(b, y, B')$, together with the law of motion of aggregate debt $B' = \Gamma(B, y)$. Specifically, we accomplish this using the following steps.

We guess an initial law of motion of aggregate debt $\Gamma^0(B, y)$. We then solve for the optimal value functions v^R and v^D and the optimal debt policy b' using value function iteration for all combinations of $(b, y; B')$. We update the law of motion of aggregate debt $\Gamma^1(B, y)$ such that $b'(b, y, B') = B'$. We iterate these procedures until $\Gamma(B, y)$ converges. If there exist more than one fixed point, we take the B' that gives the smallest debt.

We then update the default decision $d^1(B, y)$ by solving the government's problem in equation (1). Accordingly, we update the bond price schedule. In order to minimize spurious movements in the bond price, we interpolate the bond price schedule. To do so, we first interpolate the value functions v^R and v^D over the income shock y. We next find an income level $\hat{y}(B)$ at which $v^R(B, \hat{y}(B), B') = v^D(\hat{y}(B))$ for each B. Note that $\hat{y}(B)$ is not restricted to the discrete shock levels. We then update $q^1(B', y) = (1 - \int_{-\infty}^{\hat{y}(B)} f(y'|y)dy')/(1 + r)$. We iterate over the bond price schedule until it converges.

Two computational issues are worth mentioning. First, we use the discrete state-space technique with 30 endowment grid points and 1600 asset grid points. Hatchondo et al. (2010) show that the discrete state-space technique is likely to introduce spurious interest rate movements if the grid points of the state space is too coarse. Given the complexity of our model, we assume that the Markov endowment process is exogenously given by a 30-state Markov chain, obtained using a quadrature based method of Tauchen and Hussey (1991).[23] To check whether our asset grids are fine enough to offer robust results, we increase the number of asset grids to 2000 and find that the results remain almost unchanged.

Second, the fixed point mapping for the aggregate borrowing function has more than one fixed point. The existence of the fixed point in the aggregate borrowing function is straightforward to establish. Consider an aggregate borrowing function which specifies aggregate borrowing for all (B, y) to be so large that the government will default for sure next period. The bond price for such borrowing is zero, and the individual household is indifferent between all levels of debt. Without loss of generality, we can set individual borrowing the same as aggregate borrowing. Thus, we have established the existence of the fixed point, though this particular fixed point is not interesting because there is no borrowing and lending in equilibrium.

To find an "interesting" fixed point, we need to start with the initial guess for the bond price to be the inverse of the risk free rate instead of zero. To illustrate the fixed point mapping, we have plotted the debt choice of households $b'(B, y; B')$ over aggregate debt choices B' for different levels of aggregate debt and income (B, y) in Figure 11. The equilibrium debt choice b'^* is given by the intersection of the function b' and the 45-degree line. In general, as aggregate borrowing increases, the bond price declines and individual borrowing decreases. When aggregate debt is large enough such that the bond price is zero, households are indifferent between all debt levels. Thus, the individual debt choice becomes a correspondence instead of a function for this region of B', which produces a continuum of fixed points. In the case where $q(B', y)$ equals zero, we can set individual borrowing b' to zero without loss of generality, which implies a unique "interesting" fixed point. Or equivalently,

[23] The range of our discretized endowment shocks is larger than the range of the Argentina output process. In addition, the simulated income series from our discrete shock process captures pretty well various moments of the data.

we can select the fixed point with smallest debt.

Figure 11: Fixed Point Mapping

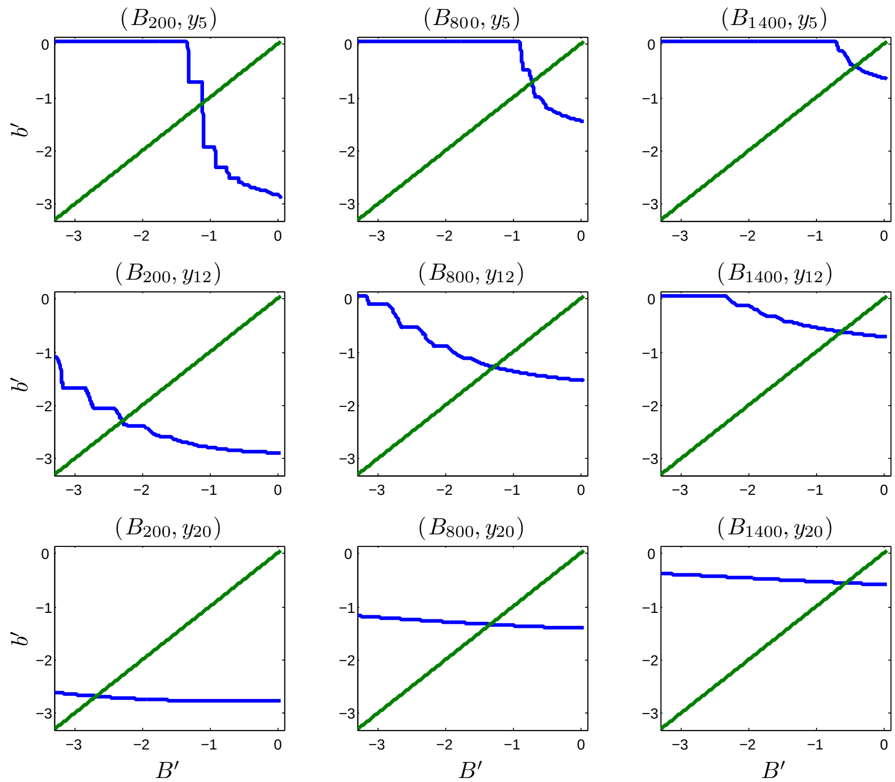

The computation algorithm for the model with centralized borrowing is simpler. We discretize the state space (B, y). We start with a guess for the bond price schedule, $q_C^0(B', y) = \frac{1}{1+r}$ for all (B', y). We next solve the optimal value functions W^R, W^D and the optimal policy function B' using value function iteration. We then update the default decision based on W^R and W^D. We finally update the bond price schedule using a smoothing method analogous to that described above. We repeat the above procedures until the bond price schedule converges.

Our computation strategy is different from Hatchondo et al. (2010) in solving the optimal debt decision. Like most studies in the literature, we use the grid search method over the discretized space. By contrast, Hatchondo et al. (2010) interpolate the value functions over the asset space and use the first-order condition to solve for the optimal debt decision. Hatchondo et al. (2010) show that when the number of the grid points are large enough, the two methods give very similar solutions. On the other hand, both their and our papers smooth the bond price function by interpolating the value functions over endowment and finding the cutoff endowment level that makes the government indifferent between repaying and defaulting. To see the effect of the bond price interpolation, Figure 12 shows the bond price schedule before and after the interpolation for the model with centralized borrowing. The discrete state-space (DSS) technique causes discrete jumps in the bond price, while the interpolation method removes spurious movements in the bond price. Simulation results show that the interpolation greatly reduces the volatility and countercyclicality of the interest rate spread. Also, it reduces the default rate substantially, which suggests that the DSS method might overestimate the default likelihood.

Figure 12: Bond Price Schedule with and without Interpolation

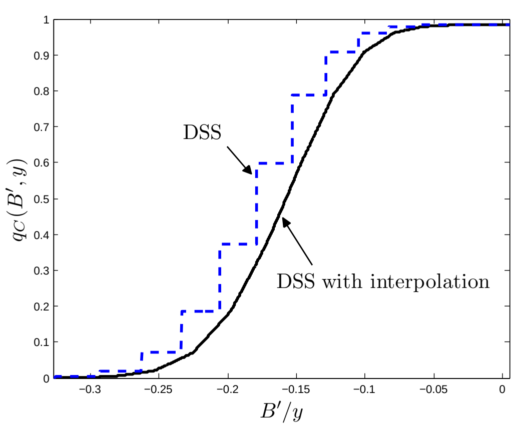

Note: The displayed bond price schedules are for the median income shock.

Appendix B – Sensitivity Analysis on Default Penalties

In this appendix, we report the simulation results for different default penalties. Table 2 contrasts the relevant statistics of the two models for different parameter values of the asymmetric default penalty specification. The decentralized borrowing model generates overborrowing when default penalties are lenient (low λ and high θ), and underborrowing when default penalties are severe (high λ and low θ). Moreover, regardless of default penalties, the decentralized borrowing model generates substantially higher mean spreads, larger default rates and lower welfare than the centralized borrowing model.

Table 2: Varying Default Penalties: Asymmetric Income Loss

	λ				θ			
	0.06	0.10	0.12	0.16	0.05	0.15	0.25	0.35
				Decentralized				
mean(B/y)	−12.73	−23.41	−27.86	−32.02	−30.69	−16.45	−10.57	−7.82
mean(spread)	10.34	11.42	11.30	7.62	9.17	11.15	10.55	10.54
prob(default)	3.37	3.10	2.67	0.95	2.17	3.41	3.37	3.52
welfare	9.90	9.86	9.82	9.79	9.81	9.88	9.91	9.88
				Centralized				
mean(B/y)	−10.99	−22.25	−27.69	−32.20	−32.32	−12.88	−7.10	−4.88
mean(spread)	0.41	0.38	0.32	0.07	0.03	0.56	0.90	1.19
prob(default)	0.11	0.14	0.07	0.01	0.01	0.15	0.30	0.36
welfare	9.95	9.92	9.91	9.90	9.91	9.93	9.94	9.95

Table 3 shows the results for different parameter values of the symmetric default penalty specification. Different from the case with the asymmetric default penalty, the decentralized borrowing model consistently generates lower equilibrium debt independent of the default penalty parameter values. Similar as in the asymmetric default penalty case, decentralized borrowing generates higher mean spreads, higher default rates, and lower welfare.

Appendix C – Sensitivity on Restricted Simulation Samples

This appendix conducts sensitivity analysis on restricted simulation samples. Specifically, we compute the model statistics based on the simulation episodes in which default ends and starts 70–78 periods apart. The results are reported in Table 4 below. By contrast, the benchmark model statistics are computed on the 74 periods that are followed by default. The key difference is that the restricted simulation episodes start with zero debt, while the benchmark simulations do not necessarily start with zero debt. In the model with decentralized borrowing, the restricted sample produces a larger mean spread (12.20% versus 11.25%) and a lower debt to output ratio (19.34% versus 22.48%). Similarly, in the recalibrated model with centralized borrowing, the restricted sample produces a larger mean spread (8.96% versus 7.30%) and a lower debt to output ratio (4.87% versus 7.23%). On the other hand, the key patterns between the centralized and decentralized borrowing models are unchanged in the restricted simulation sample.

Table 3: Varying Default Penalties: Symmetric Income Loss

	η				θ			
	0.005	0.01	0.02	0.03	0.05	0.1	0.15	0.35
				Decentralized				
mean(B/y)	−2.91	−5.92	−11.53	−16.59	−17.41	−11.53	−8.61	−4.45
mean(spread)	10.89	7.20	5.66	4.29	6.86	5.66	5.85	7.45
prob(default)	4.21	2.41	2.23	1.29	2.48	2.23	2.43	2.79
welfare	9.88	9.86	9.82	9.84	9.86	9.82	9.82	9.78
				Centralized				
mean(B/y)	−3.99	−8.20	−17.00	−26.22	−30.25	−17.00	−11.74	−5.17
mean(spread)	0.02	0.02	0.01	0.01	0.01	0.01	0.02	0.02
prob(default)	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.01
welfare	9.96	9.96	9.94	9.92	9.92	9.94	9.95	10.01

Table 4: Defaults End/Start 70–78 Periods Apart

Statistics	Data	Model	
		Decentralized	Centralized 2
mean(spread)	10.31	12.20	8.96
std(spread)	5.60	24.27	8.42
mean(B/y)	−43.36	−19.34	−4.87
std(y)	8.12	5.59	5.90
std(C)	9.47	6.42	6.38
std(TB/y)	1.75	1.90	1.51
corr(c, y)	0.98	0.95	0.97
corr(TB, y)	−0.59	−0.27	−0.19
corr(spread, y)	−0.89	−0.51	−0.56
corr(spread, c)	−0.91	−0.69	−0.62
corr(spread, TB/y)	0.70	0.79	0.46
drop in y upon default	−14.21	−12.38	−12.23
drop in c upon default	−16.01	−11.87	−12.05
discount factor β		0.97	0.93
output loss λ		0.10	0.10
re-entry probability θ		0.10	0.65
Number of episodes		118	105

CHAPTER II

The Impact of Foreign Liabilities on Small Firms: Firm-Level Evidence from the Korean Crisis

with Linda L. Tesar and Jing Zhang

2.1 Introduction

The sequence of events experienced by an emerging market undergoing a financial crisis is now all-too-familiar. Rapid economic growth and financial market liberalization encourage capital inflow, contributing to an overvalued exchange rate and increased reliance on foreign credit, usually denominated in US dollars. As economic growth and exports slow, the economy tips into crisis. The exchange rate collapses, capital flow reverses and firms find themselves unable to meet their debt requirements. Firms, and in some cases governments, become insolvent. Those deemed "too big to fail" may receive bailouts; others slash employment, declare bankruptcy or are sold to foreign owners.

While the general anatomy of crises has been well-documented,[24] the exact channels through which a financial crisis translates into a real economic contraction at the microeconomic level are less well understood. Traditional macroeconomic models predict that a depreciation of the exchange rate will be expansionary by making exports more competitive. However, if the depreciation occurs when firms are holding significant foreign-currency denominated liabilities, a negative balance-sheet effect may outweigh the export-expansion effect (Krugman 1999, Céspedes, Chang and Velasco

[24] See for example Corsetti et. al. (1999)

2004, and Feldstein 1999). In general, the literature has found ample evidence of the expansionary export effect but limited evidence of the balance-sheet effect.[25] Due to data limitations, most existing firm-level studies examine only publicly-listed firms that survived the financial crisis, and leave the extensive margin of the balance sheet effect unexplored.

In this paper, we use a detailed database on over 4000 Korean firms-both privately-held and publicly-listed-to study the impact of the 1997-1998 Korean financial crisis on firm performance. The database contains information on firms' export status, holdings of foreign debt, and total indebtedness along with a host of other firm-level characteristics. The database also provides information about firm exit during the crisis. We exploit the heterogeneity across firms to see which factors-firm size, industry, export status, exposure to foreign debt, and term structure of debt-are critical for explaining firm performance and firm exit leading up to and during the Korean crisis.

Our analysis yields three key findings. First, we find evidence of a significantly negative balance-sheet effect for small firms conditional on survival. Specifically, for a firm at the 10th percentile of the size distribution (with size measured by real assets), a one percent increase in its short-term foreign debt ratio prior to the crisis is associated with a 0.31 percent lower rate of sales growth during the crisis. Most existing firm-level studies focus on publicly-listed firms and often find either no positive balance sheet effect or a positive balance sheet effect. The reason we are able to find evidence of a significantly negative balance-sheet effect is because of the broad coverage of our dataset, which includes both well-established, publicly-listed firms as well as small, privately-held firms. This balance-sheet effect becomes insignificant when one focuses only on publicly-listed firms as in previous studies. Publicly-listed firms tend to be

[25] Benavente et. al. (2003), Bleakley and Cowan (2008), Bonomo et.al. (2003), Forbes (2002), and Luengnaruenitchai (2003) find either a positive balance sheet effect or no balance sheet effect. In contrast, Aguiar (2005), Carranza et. al. (2003), Echeverrya et. al. (2003), Gilchrist and Sim (2007) and Pratapa et.al. (2003) find some evidence of a negative balance sheet effect.

larger, are more likely to be exporting firms and are more likely to survive the crisis than an average Korean firm.

Second, we find strong evidence of the balance sheet effect on small firms at the extensive margin: foreign debt holdings are a significant predicator for firm exit during financial crisis. Consider a percentage point increase in the pre-crisis foreign debt ratio for firms at varying sizes.[26] The marginal impact on the probability of exit increases from about zero for a firm of median size to 0.28 percentage points for firms at the 10th percentile. Again, the impact of the crisis on firm exit is missed in samples that focus on publicly-listed firms that survived the 1997-98 contraction. Our dataset makes it possible to study firm exit during the Korean crisis, which accounts for nearly 20 percent of the decline in aggregate sales in the peak year of the crisis. Analysis of exit rates underscores the devastating impact of the crisis on small firms that had foreign liabilities prior to the crisis.

Third, we find a strong export-expansion effect: exporters experience smaller declines in sales growth during the crisis than non-exporters. While in principle exports provide a natural hedge against the negative effects of an exchange rate depreciation, many firms do not export and therefore do not benefit from this channel. Our data suggests that about 70 percent of Korean firms that carried foreign currency debt on their balance sheets at the time of the crisis were not engaged in exporting. Moreover, for the smallest quartile of firms, 90 percent of firms with foreign debt holdings were non-exporters. Therefore, a significant fraction of the population of Korean firms-importantly, many small firms-entered the crisis with exposure to balance sheet risk with no offsetting benefits of an improvement in global competitiveness.

Paradoxically, we find that large firms with more exposure to foreign debt experience smaller declines in sales growth during the crisis. Similar results have been documented by many studies in the literature focusing on large and publicly-listed

[26]Specifically, the ratio of foreign debt holdings to liabilities is set at the mean level conditional on having foreign debt, and all other characteristics are set at the sample mean.

firms. We also find that large firms with more exposure to foreign debt are less likely to exit during the crisis. The exact interpretation of these findings is unclear. We suspect that omitted variable bias may be the reason behind these findings. Large firms, like publicly-listed firms, are more likely to hedge exchange rate risk and have access to other means of financing during the crisis. Omitted variable bias may be less severe for small firms than for large firms, so we are able to find evidence of the balance-sheet effect for small firms.

We perform counterfactual exercises to illustrate the importance of these various channels (the balance-sheet effect and the export-expansion effect), taking into account both the contraction in sales as well as firm exit. These exercises demonstrate the importance of firm heterogeneity in assessing the role of foreign debt in the crisis.[27] For large firms, an increase in the foreign debt ratio has very little impact, if any, on firm performance. Similarly, the predicted exit rate changes very little as large firms tend to be hedged through exports. For the bottom quarter of firms, however, an across the board increase in foreign debt predicts a 1.6 percentage point decline in total sales growth conditional on survival. Their predicted exit rate increases by 7.4 percentage points, and the overall predicted decline in sales growth-taking into account sales contraction and firm exit-is nearly 8 percentage points. About 80 percent of the decline in sales growth is explained by firm exit for these small firms. What these experiments suggest is that in assessing exposure to exchange rate risk, it is important to know which firms are carrying foreign currency liabilities and whether those firms are also exporting firms.

The paper is organized as follows. Section 2.1 briefly presents macroeconomic dynamics of the Korean financial crisis. Section 2.2 describes the dataset. Section 2.3 focuses on the surviving firms and presents the evidence of the balance-sheet effect from the cross-sectional regression analysis. Section 2.4 focuses on the exit margin

[27] In this counterfactual experiment, we increase firm leverage by ten percentage points and assume that all of the increase is in short-term foreign debt.

and documents the balance-sheet effect on firm exit during financial crisis. We also conduct counterfactual experiments in this section. Section 2.5 concludes.

2.2 Macroeconomic Dynamics of the Korean Financial Crisis

In the years preceding the Asian financial crisis, South Korea was one of the fastest growing economies in the world, with sustained high growth rates for more than two decades. Beginning in late 1997, the Korean economy entered a severe economic contraction. Some indicators of the magnitude of the crisis are illustrated in Figure 13a, which shows real GDP, consumption, investment and total employment normalized to their 1997 values.[28] The declines were big: from peak to trough real GDP declined 7%, real consumption fell 14%, real investment fell 35%, and employment dropped 5%. During the crisis, the current account displayed a sudden reversal of over 15 percentage points, shifting from a negative balance of 4% of GDP to a positive 12% of GDP (Figure 13b). While the crisis was deep, it was also mercifully brief. By 1999 real GDP and consumption returned to levels above their pre-crisis values.

During the boom years, Korean firms and households dramatically increased their reliance on credit. Between 1995 and 1997, total private credit as a share of GDP increased from 104 percent of GDP to almost 120 percent of GDP (see Figure 13d). Much of the credit expansion took the form of borrowing from abroad. Figure 13c shows that external debt peaked in 1997 at 60 percent of GDP, with over a third of total borrowing with maturities of one year or less. The declines in both total private credit and external debt as shares of GDP in 1997 to 2000 illustrate the dramatic deleveraging that occurred in Korea in the aftermath of the crisis.

Figures 13e and 13f show the dynamics of two key prices: the nominal exchange rate (Korean won relative to the US dollar) and the nominal interest rate (the monthly money market rate). As shown in Figure 13e the nominal exchange rate depreciated

[28] The plot shows annualized data-the crisis hit in the fourth quarter of 1997.

by almost 100% during the last weeks of 1997, peaking in January 1998. Thereafter, the won fully floated against the dollar. It appears that there was significant overshooting of the Korean won-between late-1997 and mid-1998 the won appreciated relative to the dollar although it did not return to its pre-crisis level. The short-term interest rate (Figure 13f) also shot up during the crisis, increasing from its pre-crisis range of 10-15 percent to a peak of 25.6 percent in January 1998.

The severity of the crisis has been attributed to high rates of leverage in the economy, particularly in the form of external debt, coupled with a sudden, unanticipated (and therefore unhedged) exchange rate depreciation. Despite the large literature on this topic, there has been little microeconomic evidence to support the connection between financial variables and real economic activity. We next describe the firm-level data that we will use to analyze the linkages between balance sheet risk and firm performance.

2.3 Description of Firm-Level Data

We obtain firm-level data from the Korea Information Service, Inc. (KIS), a provider of financial and corporate data for Korean firms. The underlying source of the data is the annual financial statements of all Korean firms with assets over 7 billion won.[29] The KIS removes liquidated firms from the dataset, and therefore the main dataset contains only surviving firms. We obtained additional information on liquidated firms from the KIS in a secondary database.[30]

The KIS data have several advantages over the data that have been employed in earlier studies of financial crises in emerging markets.[31] First, the KIS data include

[29] Firms with assets of 7 billion won or more are required by the Act on External Audit of Joint-Stock Corporations to report audited financial statements to the Financial Supervisory Commission, which is then compiled by the KIS. Some firms with assets less than 7 billion won voluntarily report their financial statements and show up in the dataset.

[30] Data on liquidated firms were available by special request from the KIS.

[31] Examples, among many others, include Aguiar (2005), Bleakley and Cowan (2008), Borensztein and Lee (2002), Forbes (2002), Gilchrist and Sim (2007), Kalemli-Ozcan et al (2009), and Martinez

firms that are not listed on the Korean stock exchange. The KIS data reveal that publicly-listed firms are only a fraction of the population of Korean firms and they provide a skewed portrait of the impact of the crisis at the micro-level. The KIS data also provide information on foreign currency denominated debt versus domestic debt[32] as well as the maturity structure of the debt. The database contains firm-level information on whether a firm is an exporter or not, allowing us to disentangle the export-expansion effect from the balance-sheet effect of an exchange rate depreciation. Finally, the merged database allows us to study firm exit, a margin of adjustment during the Korean crisis that has not heretofore been studied.

2.3.1 Characteristics of Surviving Firms

Table 5 provides summary statistics for surviving firms. We focus on the 1994-1999 sample period to capture the effects of the financial crisis. We exclude firms in the financial sector. In order to limit the influence of outliers, we eliminate observations in the top and bottom 1 percent of the sample in terms of the sales growth rate and the profit rate. When firms are sorted by industry, about 62 percent of firms are in the manufacturing sector, 14 percent in wholesale, retail trade and transportation, 11 percent in construction and utility, and another 13 percent provide other services. These industry shares are fairly constant over the 1994-1999 period. The full sample of firms, shown in line 1 of the top panel, starts with a sample size of 3,151 and increases over the 1994-1999 period. The increase in the sample size over time is not surprising given that the cutoff for coverage (7 billion won) is fixed in nominal terms; as the economy grows and there is inflation, the number of firms above this cutoff

---

and Werner (2002). All these papers focus on publicly listed firms. Bleakley and Cowan (2008) have no information on export status and Forbes (2002) uses total debt statistics instead of foreign debt. All these papers, except Kalemli-Ozcan et al (2009), have no exit information. In Kalemli-Ozcan et al (2009) firms rarely exit, so the extensive margin plays a limited role in their study.

[32]The KIS does not provide the currency denomination of foreign debt. However, other sources indicate that the majority of foreign borrowing was denominated in US dollars. According to Kwon (2005), prior to the crisis 96 percent of foreign debt of publicly-listed firms was in US dollars, 3 percent in yen, and 1 percent in other currencies.

will obviously increase.

The mean age of firms (Table 5, line 2) is 15 to 17 years. In the first year, the median level of total assets (line 4) is 20 billion won, about triple the cutoff level for inclusion in the database. The mean level of real assets (line 3) is dramatically larger at 112 billion won, suggesting that the full sample covers many smaller firms. As we show below, inclusion of relatively small firms is critical for identifying the balance-sheet effect on firm performance during the crisis.

The focus of our analysis will be firm performance during the crisis as measured by sales growth rates.[33] Annual real sales growth rates and profit rates are shown in lines 5 and 6 of the top panel of Table 5. The median real sales growth rate is in the 10-15 percent range in the pre-crisis period. The crisis occurred in late 1997, and median real sales growth drops off to 6.7 percent that year and then plummets to 10.3 percent in 1998. The profit rate is defined as the ratio of the pre-tax profit and the previous-year sales. The median profit rate is around 3 percent in the pre-crisis years, and drops to 2 percent in 1997 and 2.3 percent in 1998.

Firm-level financial statistics are shown in lines 7 to 10. The leverage ratio (line 7) is defined as total liabilities over total assets. The short-term debt ratio (line 8) is the share of short-term debt in total liabilities. The foreign debt ratio (line 9) is computed as the ratio of foreign debt to total liabilities. The short-term foreign debt ratio (line 10) is the share of short-term foreign debt in total liabilities. The mean leverage ratio declines after the crisis from 76% in 1994-1997 to 67% in 1999. The short-term debt ratio is relatively constant over the period of 1994-1999 at around 30%. The foreign debt ratio is about 4% before the crisis and rises to 6% in 1997 in part due to the exchange rate depreciation. The short-term foreign debt ratio is about 1.7% before the crisis and rises to 2.2% in 1997. The number of firms with foreign debt exposure (line 13) is large: about 40 percent of the full sample of firms

[33] We also studied alternative performance measures of the pre-tax profit/sales ratio and the investment/capital ratio. The results are generally similar and are reported in the appendix.

carried foreign-currency denominated debt on their balance sheets in 1996. For those firms reporting foreign liabilities, the average foreign debt ratio was 12 percent in 1996.

Figure 14 compares the level of foreign currency debt of the banking sector and the sum of foreign currency debt of the firms in our sample. In both cases debt is decomposed into short-term and long-term debt, where short-term debt is defined as debt with original maturity of one year or less. External debt of both banks and private firms in our sample increased in the years preceding the crisis, with short-term debt accounting for roughly half of all external liabilities. This pattern is not surprising because the majority of foreign debt holdings by Korean firms are channeled through the domestic banking sector.

Previous analyses of emerging market crises suggest that exports may have provided firms with a natural hedge for foreign currency exposure–a depreciating currency will increase the cost of dollar–denominated debt service, but will increase the firm's competitiveness in foreign markets.[34] Firm exports as a share of total sales are reported in line 11 of the top panel of Table 5. The mean export to sales ratio is around 6 percent in our sample period. The fraction of exporting firms (line 12) in the full sample ranges from 13 to 20 percent. Conditional on exporting, the average export to sales ratio is around 30 percent.

The bottom panel of Table 5 reports summary statistics for publicly-listed firms, which ac-count for 20 to 25 percent of the full sample. Publicly-listed firms are older, bigger and more profitable than an average firm in the full sample. They tend to have a smaller decline in sales growth during the financial crisis. They have somewhat lower leverage ratios and short-term debt ratios. They are also more exposed to foreign-currency denominated debt and are more likely to be exporters. Firms holding foreign-currency denominated debt constitute about 64 percent of the sample

[34]See Aguiar (2005) for Mexico, Bleakley and Cowan (2008) for five Latin American countries, and Luengnaruemitchai (2003) for six East Asian countries.

of publicly-listed firms but only 39 percent of the full sample in 1996. The fraction of firms that are exporters is about 27 percent among publicly-listed firms, while only 16 percent among the full sample in 1996. Conditional on having dollar debt, the mean foreign debt ratio and the export-sales ratio are similar across these two samples.[35]

An important issue is the extent to which our sample of firms is representative of the dynamics of the aggregate economy. While our empirical work will exploit heterogeneity between firms, our results could be viewed with suspicion if our sample of firms exhibits aggregate sales behavior during the crisis that is dramatically different from the dynamics of aggregate economic activity in Korea. To address this issue, Figure 15a shows the sum of firm sales as a ratio of GDP. The top line is the sum of all sample firms relative to GDP. The ratio is just under 1 in 1994 and increases to about 1.4 in 2000 as more firms are brought into the sample.[36] The figure also shows the ratio for publicly-listed firms, which tops out at about 0.9. Figure 15b compares the time series of real GDP growth over 1994-1999 to median real sales growth for our full sample of firms. Not surprisingly, there is more variation in the sales growth, but the shape of the two curves is similar. Both series pick up the dramatic fall in economic activity in 1998 and the recovery in 1999. This suggests that the patterns we see in firm-level data are consistent with aggregate macroeconomic dynamics.

[35] Another group of Korean firms that has received a great deal of attention is the subset of firms belonging to chaebols. Chaebols are South Korean conglomerates composed of many companies clustered around one parent company. As the literature has emphasized, membership in a chaebol can provide insurance to firms through interlocking contracts and financial linkages. See Borensztein and Lee (2002), Lee et. al. (2000), and Min (2007). Our dataset includes roughly 230 firms that are part of the top 30 chaebols. Their characteristics tend to be similar to those of publicly-listed firms with several exceptions. First, the size of a chaebol firm, as measured by mean real assets, is more than twice the size of publicly-listed firms, and about seven times larger than the mean firm in the full sample. Second, the chaebols tend to have larger sales growth rates but lower profit rates than the publicly-listed firms. Third, the chaebols have much larger leverage ratios and greater exposure to foreign debt than the publicly-listed firms. Finally, the chaebols have smaller export/sales ratios than the publicly-listed firms. We include a chaebol dummy in our cross-section analysis to test for the role of network linkages on firm performance. No chaebols exited from the sample prior to the financial crisis.

[36] These numbers are smaller than the output/GDP ratio for the US economy. Based on BEA data, the ratio of gross output of all industries excluding the financial industry to GDP ranges from 1.64 to 1.7 between 1994 and 2007. Thus, firm coverage of the KIS database might be somewhat less complete than the BEA coverage.

2.3.2 Characteristics of Liquidated Firms

We now turn to liquidated firms in the sample. The KIS database provides a list of firms that submitted a notification of closing business to the court system and balance sheet information for these firms before their liquidation.[37] Table 6 provides summary statistics of firms that exited during the 1994-99 period-line numbers are identical to those in Table 5 for ease of comparison across the categories of firms. The exit rate in our sample–shown at the top of the table–was around 2 percent in the pre-crisis years, doubled to about 5 percent in 1997 and remained high at around 4 percent in 1998.[38] The exit rate dropped to 1 percent in 1999. It should be noted that no publicly-listed firms filed a notification of closing throughout the 1994-1999 period. No chaebol firms exited before the crisis, and some did exit during the crisis.

Comparing liquidated firms with all firms (recall Table 5 and note that the statistics in Table 6 are for the year preceding firm exit), we see that liquidated firms tend to be younger and much smaller in size than the average firm. Before they exit, firm-level profit rates are very low and negative.[39] Prior to exit, liquidated firms are less likely to be exporters and carry substantially more debt, particularly short-term debt, relative to the average firm. Liquidated firms are less likely to have foreign debt, and have smaller foreign debt ratios than the average firm. They also tend to be concentrated in the construction and manufacturing sectors.

In panel B of Table 6, we decompose the decline in annual aggregate firm sales growth into the drop in sales of surviving firms (the intensive margin), and the drop

[37] The list of liquidated firms does not include reorganized firms or firms that were sold to a foreign company. Thus, our exit data underestimates the severity of bankruptcy in crisis. The dataset does not allow us to precisely track entry. Firms may appear in the database either because they are newly established or because they reach the 7 billion won criterion.

[38] The exit rate in year t is computed as the number of firms that exited in year t divided by the sum of the number of surviving firms from year t-1 to t and the number of firms that exited in year t.

[39] For exiting firms, we have only after-tax profits instead of pre-tax profits. The profit rate for exiting firms is thus computed using after-tax profits.

due to firm exit (the extensive margin).[40] Consider the change in total firm sales between year t and year t+1. Some firms in year t continue in operation in year t+1, and we refer to these firms as "surviving firms." The remaining firms liquidate and exit, and we refer to them as "exiting firms."[41] The aggregate net sales growth equals the ratio of total sales of surviving firms in period t+1 and total sales of both surviving and liquidated firms in period t minus 1. We decompose aggregate sales growth into the intensive and extensive margin in the following way:

$$\frac{\text{sales}_{\text{survive},t+1}}{\text{sales}_{\text{survive},t} + \text{sales}_{\text{exit},t}} - 1 = \frac{\text{sales}_{\text{survive},t+1} - \text{sales}_{\text{survive},t}}{\text{sales}_{\text{survive},t} + \text{sales}_{\text{exit},t}} - \frac{\text{sales}_{\text{exit},t}}{\text{sales}_{\text{survive},t} + \text{sales}_{\text{exit},t}} \tag{13}$$

where the left hand side is the aggregate net sales growth, the first term on the right hand side is the intensive margin and the second term is the extensive margin. The intensive margin is the ratio of the change in total sales of surviving firms between year t+1 and t and total sales in period t, and the extensive margin is the ratio of total sales of exiting firms and total sales of all firms in year t.

As the table shows, the contribution of the extensive margin to aggregate sales growth is small prior to the crisis-about 3 percent of total sales growth in our sample. In the crisis years, however, the extensive margin becomes substantially more important, accounting for 18 percent of the fall in aggregate sales growth in 1998.

2.4 Cross-sectional Analysis of Firm Performance

Before turning to the regression analysis, we first plot the time series of median sales growth for different subgroups of firms. We restrict the sample to the firms

[40] In this analysis, we abstract from the entry margin because the KIS dataset does not cover many entering firms. We doubt that the entry margin plays an important role during the financial crisis.

[41] In this analysis, we ignore the contribution to total sales growth by firms that newly enter the database in period t+1.

that report relevant statistics throughout 1994-1999. We classify firms into different groups according to their characteristics in 1996. Figure 16 shows median sales growth for firms by industry, firm size, export status, leverage, short-term debt and foreign debt as a share of total liabilities. The overwhelming message of Figure 16 is that the economic contraction was a macroeconomic phenomenon. While there are some differences across firms–for example sales of non-exporters contracted more sharply than exporters, and sales of the construction and utility industry had the deepest fall in 1998–virtually all sectors and all types of firms experienced a deep contraction in 1998 and a sharp recovery in 1999. This suggests that to the extent differences in firm-level characteristics are important for understanding the crisis, they will only explain a fraction of the overall variation, and will likely work through interaction effects or through firm exit.

The general form for the cross-section regressions is shown in the following equation:[42]

$$\text{SALES GROWTH}_i = \alpha + \beta \text{CHAR}_{i,-2} + \varepsilon_i. \tag{14}$$

The dependent variable is firm i's annual real sales growth. We perform the regression analysis for two time periods-the crisis period (characteristics in 1996 as explanators for the sales growth rate between 1997 and 1998) and the pre-crisis period (characteristics in 1994 as explanators for the sales growth rate between 1995 and 1996).[43]

[42]Our goal is to account for the cross-sectional variation in firm performances during the crisis, and to relate this variation to firm-specific pre-crisis characteristics. An alternative would be to use a panel specification with firm fixed effects, and estimate how within-firm variation in debt holdings and export sales affects variation in firm performances over time. In that case, the impact of the crisis would be estimated through an interaction of lagged firm characteristics with the crisis dummy. We do not pursue this strategy for three reasons. First, such a specification would answer a different, much more narrow question: how does the crisis affect the relationship between debt holdings or export sales and sales growth within a firm? Second, firm fixed effects soak up explanatory power of interesting and informative time-invariant firm characteristics. Third, the short-time dimension of our dataset implies that we have limited variation to exploit.

[43]We repeat the analysis with alternative measures of firm performance: the profit rate and the investment rate. The results are reported in the appendix. The main findings are broadly similar to those we report for sales growth. To ease exposition, we will focus primarily on the results for sales growth.

In the baseline specification, firm-level characteristics include size (the log of firm real assets), age, chaebol status, leverage ratio, short-term debt ratio, export/sales ratio, and foreign debt ratio. All variables are in real Korean won. We include a two-digit industry dummy to control for industry-specific effects. In the second specification, we also include interaction effects between firm size with the foreign debt ratio, the leverage ratio, and the short-term debt ratio to allow these variables' effects on sales growth to differ by firm size. In the third specification, we decompose the foreign debt ratio by maturity to examine whether firms with varying foreign debt maturities have differential firm performance.

2.4.1 Cross-Section Results for Publicly-Listed Firms

Table 7 shows the results for the sample of publicly-listed firms-the firms that have been carefully studied in previous analyses. Specifically, columns 1, 2 and 3 report the results for the crisis period, and columns 4, 5 and 6 report the results for the pre-crisis period. Our results confirm the results generally reported in the literature. In the pre-crisis period, chaebol status is positively related to firm performance while firm age is negatively related to firm performance. However, there is no significant effect of leverage, exports or balance-sheet variables. The results for the crisis period are somewhat different. Though the chaebol status continues to be positively related to firm growth, the age effect disappears. More importantly, the export-sales ratio now appears with a positive, statistically significant coefficient, confirming the export-expansion effect found in previous studies.

In terms of the balance-sheet effect, the first specification, in which financial variables are not interacted with firm size, presents a puzzling result. The coefficients on the leverage ratio and on the foreign debt ratio are significantly positive. This seems to suggest that firms entering the crisis with higher leverage ratios or higher foreign debt ratios had better performance during the crisis. The literature reports similar findings for publicly-listed firms (see Bleakley and Cowan 2009). When financial variables are interacted with size (columns 2 and 3) the positive coefficient disappears (the first clue that something different is going on for smaller firms) and the coefficients on financial variables are no longer statistically significant. Thus, an analysis based on publicly-listed firms would either suggest a positive role for foreign debt (if the size effect were omitted) or no role for financial variables in explaining firm performance.

The positive coefficient on the foreign debt ratio in column 1 might be suggestive of potential omitted variable bias. If some firm characteristics that are positively correlated with a firm's ability to raise foreign debt and with its sales growth are omitted from our regression, the estimated coefficient on the foreign debt ratio would be biased upward. For example, publicly-listed firms with more foreign debt may better hedge against exchange rate risk through holding financial derivatives or foreign currency denominated assets. Their hedging decisions in turn might lead to smaller declines in sales growth during an exchange rate depreciation. Another example is that publicly-listed firms with more foreign debt may also have greater access to other forms of credit in crisis, so they experience smaller declines in sales growth. Our dataset does not include the information on holdings of financial derivatives and foreign currency denominated assets and on accessibility to financing to control for these potential sources of bias.[44]

2.4.2 Cross-Section Results for the Full Sample

Table 8 repeats the analysis for a balanced sample of firms that includes small, privately-held firms. In the pre-crisis cross-section regression, we take the sample of firms in 1994 and hold that sample fixed through 1996. For the crisis cross-section regression, we take the sample of firms in 1996, holding the sample fixed through

[44] Korean firms are required to disclose information about financial derivatives by law only after 2000.

1998. Note that this sample is about four times the size of the sample in Table 7. This analysis will still miss the impact of firm exit, however, as we include only those firms that survive for the three-year interval (1994-1996 in the pre-crisis regression and 1996-1998 in the crisis regression).

Turning first to the pre-crisis regression results in the right panel of Table 8, we see that the effects of chaebol status and age remain significant. Specifically, chaebols and younger firms are associated with faster sales growth. Moreover, there is no evidence of an export effect, similar to the results for publicly-listed firms in the pre-crisis period. The first specification (column 4) yields no significant effect for leverage, short-term debt or foreign debt ratios prior to the crisis. When these financial ratios are interacted with firm size (columns 5 and 6), leverage has a significant effect that varies with firm size. Higher leverage ratios are associated with faster sales growth rates for small firms, but slower sales growth rates for large firms. On the other hand, greater exposure to short-term debt is associated with slower sales growth rates, though the effect of short-term debt is smaller and may be positive for larger firms. Foreign debt ratios remain insignificant in the full sample of firms in the pre-crisis years.

The results are dramatically different during the crisis (the left panel of Table 8). There is a robust relationship between exports and firm sales: the coefficient on export status is positive and strongly significant across all three specifications. The effect is also economically significant. A ten percent increase in the pre-crisis export sales ratio is associated with an increase in sales growth of approximately 2 percent during the crisis. This export-expansion effect is similar to what we find in the sample of publicly-listed firms.

The main difference across the two samples of firms is the balance-sheet effect. In contrast to the findings for publicly-listed firms, the full sample shows strong evidence of a negative balance-sheet effect on small firms. Again, if the financial variables are not interacted with size, the specification in column 1 of Table 8 yields a significantly positive coefficient on foreign debt. When we include interaction terms between financial variables and firm size in column 2, the coefficient on foreign debt ratios turns significantly negative and the coefficient on the interaction term between foreign debt and size is significantly positive. Holding all the other variables constant, a one-percent larger foreign debt ratio affects sales growth by (2.817+0.134*size) percent, which monotonically increases with firm size. The impact is negative for small firms, but positive for large firms. The critical size, below which the effect of foreign debt is negative is 21.02 in terms of log real assets and corresponds to a firm at about the bottom 2 percentile in the size distribution. Thus, for most firms, a higher foreign debt ratio is associated with a higher sales growth rate during the crisis. The negative balance-sheet effect shows up only for very small firms.

The negative balance sheet effect on small firms is more prominent through short-term foreign debt. The coefficient on short-term foreign debt in column 3 is significant and large: 4.6. There is again an interaction effect with size–for large firms in the sample, the impact of short-term foreign debt is positive while for small firms the effect is negative. In this case, the cut-off point is 23.7 in terms of log real assets and corresponds to a firm at the 58th percentile in the size distribution. Thus, for firms with assets below the 58th percentile, an increase in the short-term foreign debt ratio is associated with a lower sales growth rate, all else equal. The effects are economically significant. Consider a firm with assets at the 10th percentile (log real asset of 22.1). A one percent increase in the short-term foreign debt ratio prior to the crisis is associated with a 0.31 percent lower rate of sales growth during the crisis. Note that the corresponding coefficients are similar for the sample of publicly-listed firms but were not statistically significant.[45]

[45]We included a dummy for foreign ownership to test the hypothesis that firms controlled by foreign owners have access to other credit channels and may have been buffered from the effects of the Korean crisis (see, for example, Kalemli-Ozcan et. al. 2009). Foreign-owned firms did not display different results from the full set of firms.

There are at least two possible explanations for why significant balance-sheet effects emerge in the full sample but not in the smaller sample of publicly-listed firms. One reason may be that the larger number of observations and the greater variation in the full sample yields more explanatory power. The other reason may be that omitted variable bias is more severe in the sample of publicly-listed firms than the full sample. Our conjecture is that publicly-listed firms are more likely to hedge exchange rate risk and have access to other means of financing during the crisis, than other firms in the full sample. Our data does not provide enough information for us to determine the precise reasons for the difference between the two samples, but the results are suggestive that relying on publicly-listed firms will affect one's interpretation of the impact of the crisis at the micro level.

Omitted variable bias might also help understand our paradoxical finding for large firms in the full sample. Large firms tend to be publicly-listed firms, or to have similar characteristics as publicly-listed firms. We conjecture that the positive balance-sheet effect for large firms is due to omitted variable bias, although the exact interpretation of this finding is unclear. On the other hand, the possibility of potential upward bias from omitted variables might strengthen our conclusion about the negative balance sheet effect on small firms.

The cross-section results based on the full sample support the view that both the export-expansion channel and the negative balance-sheet channel played a role during the crisis, with a particular role for exposure to short-term foreign debt. An interesting question is whether firms that were exposed to balance sheet risk were also exporters, and therefore were at least partially hedged from the negative impact of the exchange rate devaluation. Table 9 shows the decomposition of firms by export status and foreign debt holdings. The table shows that the share of non-exporters among firms that held foreign debt is 71 percent in the full sample and 66 percent in the publicly-listed sample. (The breakdown is similar for short-term foreign debt.) Thus, a significant fraction of firms that entered the crisis with foreign debt did not have a natural hedge for their currency exposure. The ratio of "non-hedged" to "hedged" firms-as measured by export status-is higher in the full sample than in the publicly-listed firms: 2.4 (=71/29) versus 1.9 (=66/34). We find that the ratio decreases with firm size; the ratio is above 8 (=89/11) for the smallest quartile and about 2 (=68/32) for the largest quartile, indicating that small firms with foreign debt holdings were more exposed to exchange rate risk.

2.5 Firm Exit During the Financial Crisis

The cross-section results pertain to firms that survived the crisis. We now perform an analysis of the factors that predict a firm's liquidation before and during the crisis. We find that foreign debt holdings are a significant predictor of firm exit, in particular for small firms, during the crisis. We then combine the intensive and extensive margin to examine the differential impact of foreign debt on firm performance by firm size.

2.5.1 Predicting Firm Exit

We run the following nonlinear probability regression on the panel of both surviving and exiting firms for the pre-crisis and crisis period:

$$P(\text{EXIT}_i = 1) = \Phi(\alpha + \theta \text{CHAR}_{i,-1}), \tag{15}$$

where P denotes the probability, EXIT is an indicator function of firm liquidation, and denotes the logistic function. In the crisis period, the dependent variable is 1 if the firm exited in 1997 or 1998, and 0 otherwise. The independent variables are firm-specific observations in 1996, to capture the pre-crisis characteristics of the firm. In the pre-crisis period, the dependent variable is 1 if the firm exited in 1995 or 1996, and 0 otherwise. Firm characteristics on the right hand side are measured in

1994. Firm characteristics include chaebol status, age, one-digit industry dummy, size, export/sales ratios, profit/assets ratios, leverage ratios, short-term debt ratios, and (short-term and long-term) foreign debt ratios. Comparing results before and during the crisis shows whether the factors that are correlated with the likelihood of firm exit during the crisis are different from those before the crisis.

The coefficients of the logit regressions are reported in Table 10. Turning first to the pre-crisis period (columns 4, 5 and 6) we see that relative to surviving firms, exiting firms tend to be younger and carry more debt, particularly short-term debt in the year preceding liquidation. Lower profits as a share of total assets significantly increase the probability of exit. Export status does not affect significantly the likelihood of exit. Turning next to the crisis period (columns 1, 2 and 3) we see weak evidence that being a chaebol member decreases the likelihood of exit.[46] Younger firms continue to have larger exit probabilities. The role of profits is less important, while leverage and short-term debt become much more important. For a nonchaebol manufacturing firm with all characteristics at the mean level, the marginal effect of a higher leverage ratio on the exit probability is six times larger during the crisis than before the crisis. The marginal effect of a higher short-term debt ratio is four times larger. The coefficient on the export sales ratio changes from positive in the pre-crisis period to negative in the crisis period, though it is still not statistically significant.

We next focus on the impact of foreign debt on firm exit. Columns 1 and 4 present a puzzling result that foreign debt does not significantly affect, if any it reduces, exit probabilities both pre-crisis and during the crisis. When interacted with firm size (columns 2, 3, 5 and 6), foreign debt has a significant effect, which varies with firm size, on the exit probabilities, suggesting that small firms with foreign debt are more likely to exit while large firms are less likely to exit. Although the coefficients on

[46] Note that in the pre-crisis period, no chaebol firms were liquidated, and therefore we cannot compare across samples. Even during the crisis, chaebol firms tended to be restructured and absorbed by other firms rather than undergo complete liquidation.

the foreign debt ratios and the interaction terms are similar pre-crisis and during the crisis, the marginal effects can be different. In a nonlinear model, the marginal effect of independent variables depends on all the covariates in the model. Especially for an interaction effect, not only the magnitude but also statistical significance varies by observation.

To examine the marginal effect of foreign debt across the two periods, the lower panel of Figure 17 plots the marginal effect of foreign debt on the exit probability (y-axis) before and during the crisis for nonchaebol manufacturing firms with different size. We fix foreign debt ratios at the mean level conditional on having positive foreign debt and all the other variables at the mean level of the sample. The solid line is the estimated marginal effect and the two dashed lines are the 95% confidence intervals. The marginal effect of foreign debt is significantly positive for small firms and significantly negative for large firms during the crisis. In contrast, the marginal effect of foreign debt pre-crisis is not generally significant except for very large firms. Thus, a larger foreign debt ratio raises exit probabilities of small firms only during the crisis.

During the crisis, for firms below the 54th percentile of the size distribution, a larger foreign debt ratio predicts a higher likelihood of exit, and for firms above, a larger foreign debt ratio lowers the probability of exit. For example, for a firm with size at the 10th percentile and all other variables at the mean level, an increase in the pre-crisis foreign debt ratio of one percentage point is expected to increase the probability of exit during the crisis by 0.28 percentage points. In contrast, if the firm is in the top decile, a one-percentage point increase in the foreign debt ratio is expected to decrease the probability of exit by 0.23 percentage points. Similar findings hold for both the short-term and the long-term foreign debt ratios.

The likelihood of exit also differs by firm size and across the two sample periods. The upper panel of Figure 17 plots the estimated likelihood of exit for the same set of firms as in the lower panel. Clearly, the predicted exit rates are higher across all firm size during the crisis than before the crisis. A firm with size at 10th percentile has a probability of exit at 7.3 percent, while a firm with size at 90th percentile has a probability of exit at 4.7 percent during the crisis. In contrast, the predicted exit rate before the crisis is only 3.5 percent for a firm with size at 10th percentile and 1.4 percent for a firm at 90th percentile.

2.5.2 Counterfactual exercise

The previous results suggest that there are important interaction effects between firm size and foreign debt, and that these effects vary across both the extensive and intensive margins. In addition, we find that the export sales are a natural hedge to foreign currency debt during the crisis. In this subsection we perform counterfactual exercises to illustrate the roles played by these various factors in accounting for the drop in firm sales during the financial crisis.

We first consider a counterfactual scenario in which each firm in our sample increases its pre-crisis leverage ratio by 10 percentage points and all of the additional borrowing is in the form of short-term foreign debt. We hold all of the other pre-crisis firm characteristics unchanged. The regression results of Column 3 of Table 8 is used to calculate the counterfactual sales growth of each firm in this scenario, conditional on survival. We then compute the average sales growth rate, weighted by 1997 sales, for each asset quartile and for the economy as a whole. Column 1 of Table 11 reports the predicted sales growth given firm characteristics as observed in 1996, and column 2 reports the predicted sales growth given the counterfactual foreign debt levels. The results illustrate the range of the impact of foreign debt on sales by firm size. Larger short-term foreign debt lowers the sales growth of small firms, but increases the sales growth of large firms. Specifically, the first (smallest) quartile sees a decline in the sales growth rate from 2.9% to 4.5%, while the fourth (largest) quartile sees an increase from 13.5% to 9%. The aggregate sales growth rate rises from 12.6% to 8.6% because large firms dominate overall sales growth.

We next turn to the extensive margin using the logit results in column 3 of Table 10. Recall that larger short-term foreign debt ratios increase the exit probability of small firms, but reduce the exit probability of large firms. The average probabilities of exit for each asset quartile given the observed characteristics and given the counterfactual short-term foreign debt ratios are reported in column 4 and 5 of Table 11, respectively. Increasing short-term foreign debt leads to a doubling of the exit probability of firms in the smallest quartile from 7.2% to 14.6%, while it reduces the exit probability of firms in the largest quartile from 6.2% to 4.5%. The overall exit rate rises from 7% to 9.3% as the foreign debt holdings increase in the economy.

We now combine the extensive and intensive margins by computing the average of the predicted sales growth rate conditional on survival and the sales growth rate of 1 conditional on exit, weighted by the survival and exit probability, respectively. See column 7 and 8 of Table 11. Incorporating both effects, we find that increasing short-term foreign debt is associated with a decline in the sales growth rate by 7.6 percentage points for firms in the smallest quartile, but is associated with a rise of the sales growth rate by 5.7 percentage points for firms in the largest quartile. This result suggests that the impact of foreign debt depends critically on what types of firms take on foreign debt. If foreign debt is concentrated in the balance sheets of large firms, which have ways to hedge against the currency depreciation in the crisis, foreign debt is not necessarily detrimental to firms' performance. On the other hand, if foreign debt is concentrated in the balance sheets of small firms, the decline in predicted sales growth is large. Note also that the extensive margin explains the majority of the decline of sales growth for most firms in the sample. For example, the extensive margin accounts for 80 percent of the decline in sales for the smallest quartile and 73 percent for the second smallest quartile. These numbers underscore the importance of taking firm exit into account when evaluating the effects of the crisis.

We conduct the second counterfactual experiment on the potentially mitigating role played by export sales during the crisis. In this scenario, we set all exporters' export/sales ratio to zero, to essentially eliminate any of the natural hedging effect of firm exports on sales growth. The predicted sales growth rates conditional on survival are reported in column 3 of Table 11. From the cross-section regression results, we know that higher export-sales ratios are associated with better firm performance. Thus, it is not surprising that the counterfactual export sales ratios lead to lower sales growth rates for all asset quartiles and especially in the largest asset quartile where most of the exporting firms appear. Overall sales growth also declines from 12.6% to 14.6%. Since the export/sales ratio does not have a large role in explaining exit probabilities, the extensive margin (reported in column 6 of Table 11) changes little from column 4. This suggests that to the extent exports provided a natural hedge for the exchange rate depreciation, they did so primarily for the largest firms in the sample and they did not shield small firms from the risk of bankruptcy during the crisis.

2.6 Conclusion

Using Korean firm-level data on both publicly-listed and privately-held firms and firm exit data, this paper finds evidence of a balance-sheet effect and an export-expansion effect. Before the crisis, firm sales growth was uncorrelated with foreign debt holdings and export sales. During the crisis, however, small firms holding more foreign debt, in particular, short-term foreign debt, experienced larger declines in sales growth. Firms with higher export sales ratios have smaller declines in sales growth during crisis. In addition, we find that small firms with short-term foreign debt are significantly more likely to go bankrupt during the crisis. The extensive margin accounts for a large fraction of small firms' adjustment during the crisis.

There are two caveats to these conclusions. The first is that the results in this paper pertain primarily to differential firm performance in the cross-section. As shown in Figure 16, most of the variation in the data is at the macro level. That is, our results can only explain whether firms with more foreign debt holdings have sharper declines in sales than firms with smaller holdings and we do not claim to provide an explanation for the overall decline in firm sales. Second, the regression analyses take firm characteristics (size, debt ratios, export status, etc.) as given in explaining next period's sales growth. Obviously, many firm characteristics are themselves choice variables, and a complete model would endogenize the full menu of firm characteristics, including firm debt, exposure to foreign currency risk and export status. We leave a more complete analysis that would address these caveats for future research.

References

Aguiar, Mark, 2005. "Investment, devaluation, and foreign currency exposure: The case of Mexico," Journal of Development Economics 78, 95-113.

Bayraktar, Nihal, Plutarchos Sakellaris and Philip Vermeulen, 2005. "Real Versus Financial Frictions to Capital Investment," European Central Bank Working Paper 566.

Benavente, Jose Miguel, Christian A. Johnson, and Felipe G. Morande, 2003. "Debt composition and balance sheet effects of exchange rate depreciations: a firm-level analysis for Chile," Emerging Markets Review 4 (2003) 368-396.

Bleakley, Hoyt and Kevin Cowan, 2008. "Corporate dollar debt and depreciations: Much ado about nothing?" The Review of Economics and Statistics 90(4), 612-626.

Bleakley, Hoyt and Kevin Cowan, 2009. "Mishmash on Mismatch? Balance-Sheet Effects and Emerging-Markets Crises," Working Paper, University of Chicago.

Bonomo, Marco, Betina Martins, and Rodrigo Pinto. 2003. "Debt composition and exchange rate balance sheet effect in Brazil: a firm level analysis," Emerging Markets Review 4 (2003) 368-396.

Borensztein, Eduardo and Jong-Wha Lee, 2002. "Financial crisis and credit crunch in Korea: evidence from firm-level data," Journal of Monetary Economics 49, 853-875.

Carranza, Luis J., Juan M. Cayo, Jose E. Galdon-Sanchez, 2003. "Exchange rate volatility and economic performance in Peru: firm level analysis," Emerging Markets Review 4 (2003) 472-496.

Céspedes, Luis J., Chang, Roberto and Andrés Velasco, 2004. "Balance sheets and exchange rate policy," The American Economic Review 94(4), 1183-93.

Corsetti, Giancarlo, Paolo Presenti and Nouriel Roubini, 1998. "What caused the Asian currency and financial crisis? Part I: Macroeconomic Overview," NBER working paper 6833.

Echeverrya, Juan Carlos, Leopoldo Fergussona, Roberto Steinerb, and Camila Aguilara, 2003. "Dollar debt in Colombian firms: are sinners punished during devaluations?" Emerging Markets Review 4 (2003) 417-449.

Feldstein, Martin, 1999. "Self-protection for emerging market economies," NBER Working Paper 6907.

Forbes, Kristine J. 2002. "How do large depreciations affect firm performance?" IMF staff papers 49.

Gilchrist, Simon and Jae W. Sim, 2007. "Investment during the Korean financial crisis: a structural econometric analysis," NBER working paper 13315.

Kalemli-Ozcan, Sebnem, Herman Kamil, and Carolina Villegas-Sanchez, 2009. "What Hinders investment in the aftermath of financial crises: balance-sheet mismatches or access to finance?" Working paper, University of Huston.

Krugman, Paul, 1999. "Balance sheets, the transfer problem, and financial crises," International Tax and Public Finance 6(4), 459-472.

Kwon, Taek Ho, 2005. "Asymmetric Exchange Rate Exposure and Foreign Currency Denominated Debt," Working paper, Chonnam National University.

Lee, Jong-Wha, Young Soo Lee and Byung-Sun Lee, 2000. "The determination of corporate debt in Korea," Asian Economic Journal 14, 333-356.

Luengnaruemitchai, Pipat, 2003. "The Asian crisis and the mystery of the missing balance sheet effect," Working paper, University of California, Berkeley.

Martnez, Lorenza and Alejandro Werner, 2002. "The exchange rate regime and the currency com-position of corporate debt: the Mexican experience," Journal of Development Economics, 69 (2002) 315-334.

Min, Byung, 2007. "Changing pattern of corporate governance and financing in the Korean Chaebols," Economic Papers 26(3), 211-230.

Pratapa, Sangeeta, Ignacio Lobatoa, and Alejandro Somuanob, 2003. "Debt composition and balance sheet effects of exchange rate volatility in Mexico: a firm level analysis," Emerging Markets Review 4 (2003) 450-471.

Appendix D

In this appendix, we report the cross-section regression results for two alternative measures of firm performance: the profit rate and the investment rate. We first describe the results on the profit rate. In the crisis regressions, the dependent variable, i.e., the profit rate, is pre-tax profits in 1998 as a share of sales in 1997. In the pre-crisis regressions, the profit rate is pre-tax profits in 1996 as a share of sales in 1995. The independent variables are the same as in the sales growth regressions. We run the regressions for both the publicly-listed firms (Table 12) and the full sample (Table 13). The results on the balance sheet effect and the export expansionary effect are very similar to the results with sales growth. Firms entering the crisis with larger foreign debt ratios have higher profit rates during the crisis (see column 1 of both tables). However, when interacting with firm size, the coefficients on the foreign debt ratio becomes negative and the coefficients on the interaction term become positive. Thus, the negative balance sheet effect shows up for small firms. This effect is not statistically significant for publicly listed firms, but it is significant at the ten percent level for the full sample and the short-term foreign debt ratio.

We then look at the results with firm performance measured by the investment rate. The investment rate is the ratio of real investment to the lagged replacement value of real capital stock. In the crisis regression, the dependent variable is firm real investment in 1998 as a share of the re-placement value of real capital stock in 1997. In the pre-crisis regression, the dependent variable is firm real investment in 1996 as a share of the replacement value of real capital in 1995. The independent variables include all the characteristics that we examined for the sales growth rate and the profit rate. In addition, we also include the lagged dependent variable as an additional regressor to pick up the persistence effect of investment.

Before introducing the results, we explain the construction of the investment rate in more details. For real investment (I_t), nominal investment is first constructed by $I_t^n = K_t^b - K_{t-1}^b + Dep_t$, where K_t^b is calculated by subtracting land and lease assets from tangible assets (all in book values from the balance sheets), and Dept is taken from the cash flow statements. Real investment is nominal investment deflated by capital goods price index. The replacement value of real capital stock (K_t) is calculated by iterating $K_t = (1-d)K_{t-1} + I_t$ backward, where I_t is real investment constructed as above and the economic depreciation rate d is assumed to be 11%, which is an average depreciation of building, structure, vehicle and machine in South Korea. The initial capital stock is measured as the real book value of capital in the year that a firm first appears the data set.[47]

We report the results for the publicly-listed firms in Table 14, and the results for the full sample in Table 15. For the publicly-listed firms, the balance sheet effect and the export expansionary effect do not show up significantly in all specifications during the crisis and before the crisis. For the full sample, we find evidence for both effects. Larger export sales ratios are statistically significantly associated with higher investment rates only during the crisis. Small firms with larger foreign debt ratios have lower investment rates both during the crisis and before the crisis. The maturity structure of foreign debt holdings does not matter for the investment rate during the crisis. Before the crisis, small firms with larger long-term foreign debt ratios have lower investment rates.

[47]We follow Bayraktar et. al. (2005) in constructing the investment rate.

Figure 13: Aggregate Data

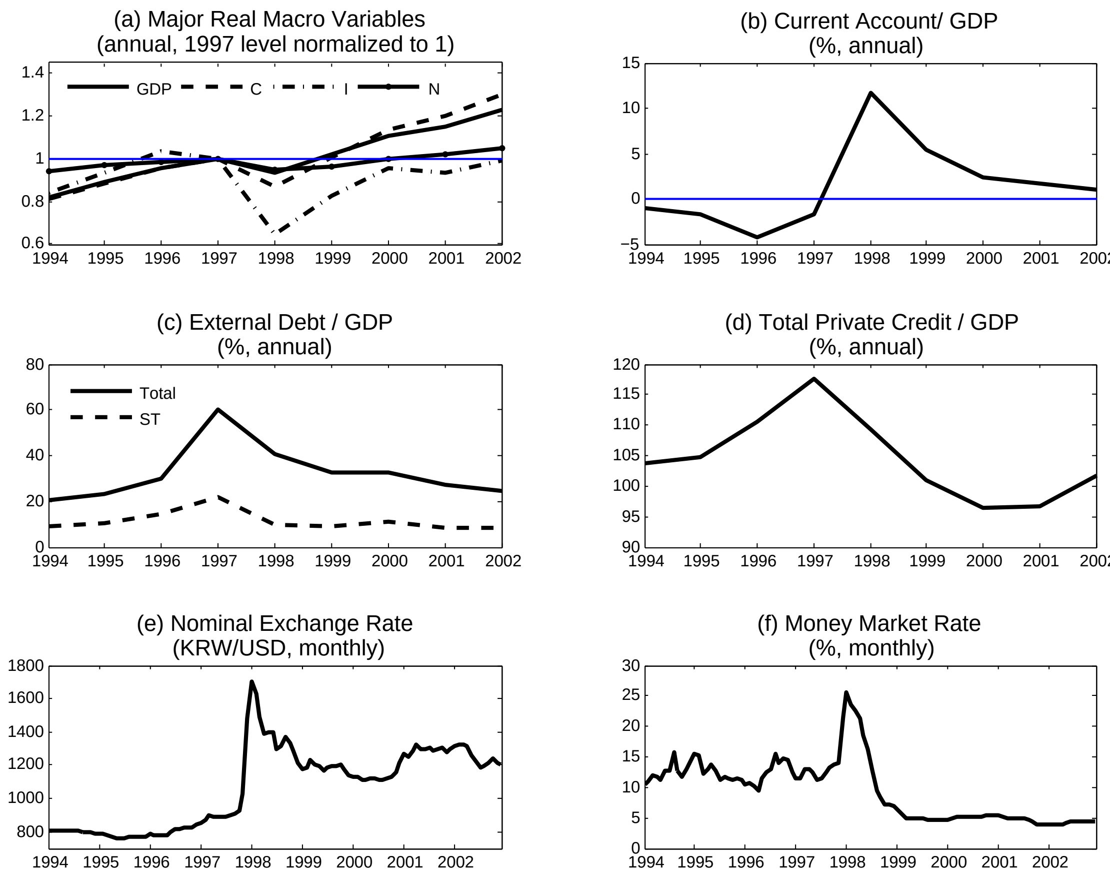

Note: The data source is Korea National Statistical Office.

Figure 14: Aggregate and Firm-Level Debt Data

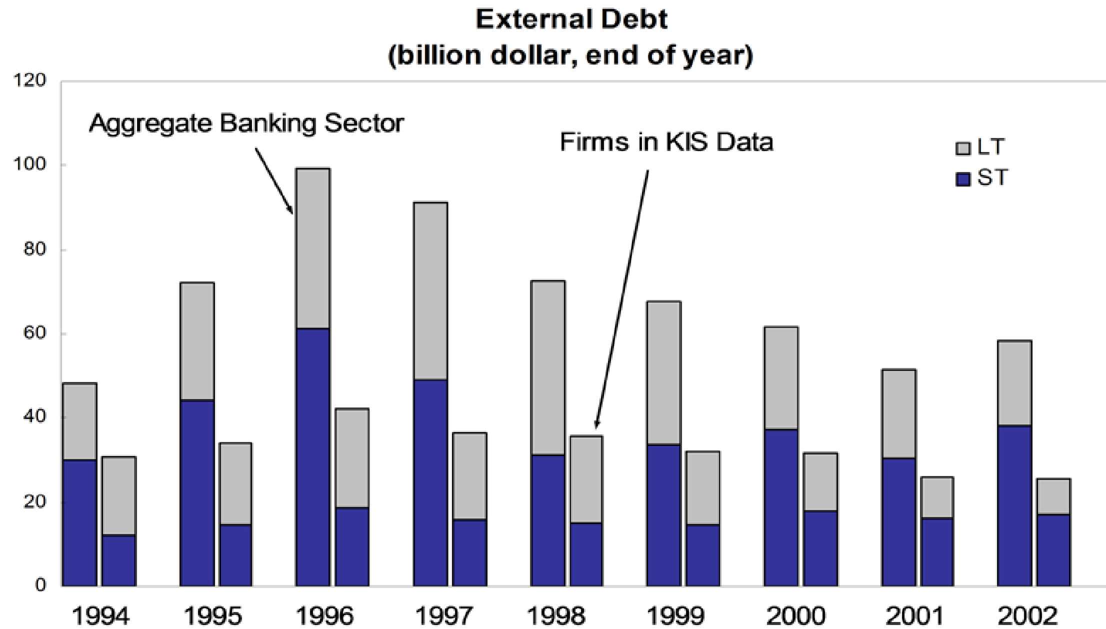

Note: Short-term debt has original maturity equal to or less than one year. The aggregate debt statistics come from Korea National Statistical Office, and the firm-level debt statistics come from the KIS-VALUE dataset.

Figure 15: Comparison of Firm Sales and GDP

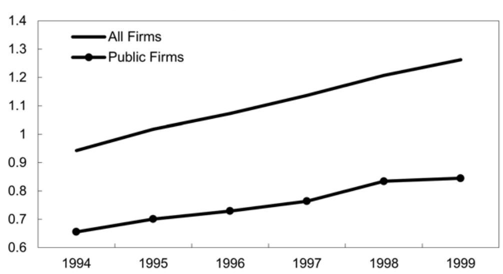

(a) Sum of Firm Sales / GDP

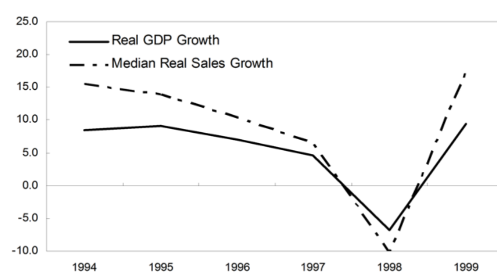

(b) Real GDP and Firm Sales Growth (%)

Note: The data sources are Korea National Statistical Office and the KIS-VALUE dataset.

Figure 16: Sales Growth of Firms with Varying Characteristics

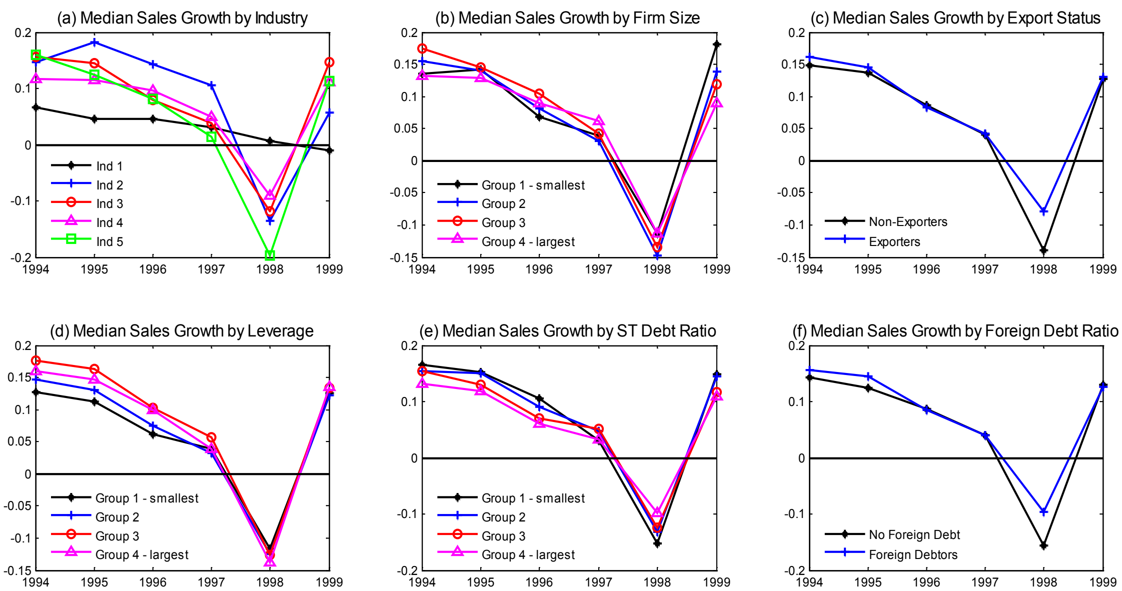

Note: Industry 1 is Agriculture, Forestry, Fishing, and Mining; Industry 2 is Construction and Utility; Industry 3 is Manufacturing; Industry 4 is Wholesale and Retail Trade and Transportation; Industry 5 is Other Services. The data source is the KIS-VALUE dataset.

Figure 17: LOGIT Regression Results

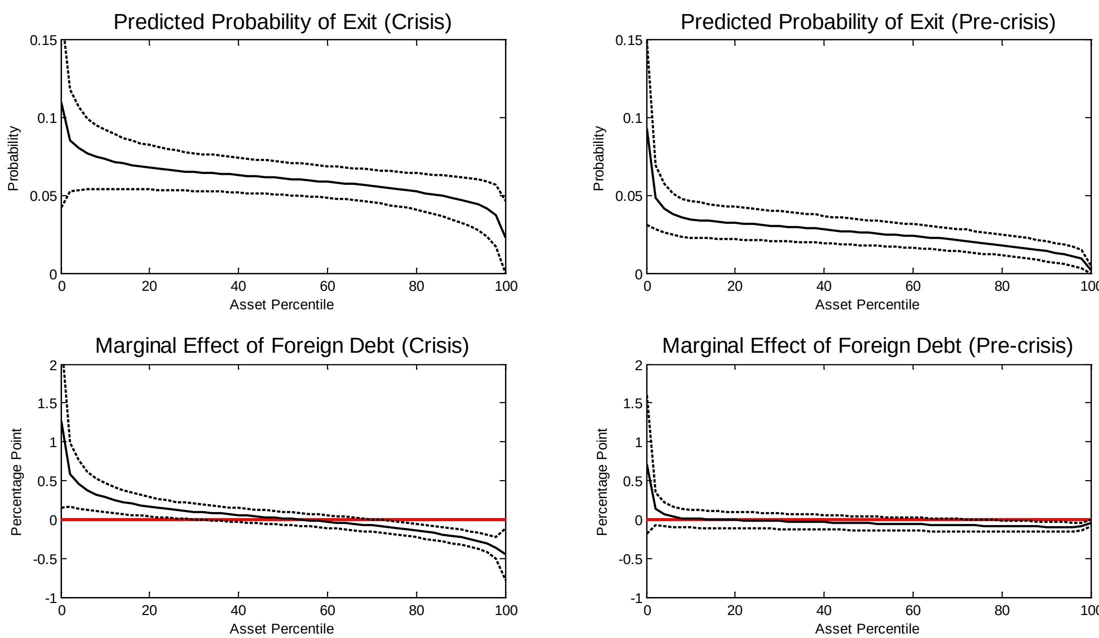

Note: In the upper panel, the solid line plots the estimated probability of exit for nonchaebol manufacturing firms with different size, the foreign debt ratio at the mean level conditional on having positive foreign debt, and all the other variables at the mean level of the sample. In the lower panel, the solid line plots the estimated marginal effect of foreign debt on the exit probability in the crisis period for the same set of firms as in the upper panel. The left panel is for the crisis period, and the right panel is for the pre-crisis period. The dashed lines are the 95% confidence intervals.

Table 5: Summary Statistics for Surviving Firms

		1994	1995	1996	1997	1998	1999
				All Firms			
1	Number of Firms	3,151	3,956	4,285	5,066	5,476	5,606
2	Mean Age	17	17	16	16	15	15
3	Mean Real Assets	112	114	109	120	95	101
4	Median Real Assets	20	19	16	15	11	11
5	Median Real Sales Growth Rate (%)	15.3	13.9	10.5	6.7	-10.3	19.6
6	Median Profit Rate (%)	3.3	2.9	2.6	2	2.3	4.7
7	Mean Leverage Ratio (%)	76	76.6	76.5	77.1	72	67.3
8	Mean ST Debt Ratio (%)	30.5	30.7	30.4	29.7	30.4	29.7
9	Mean Foreign Debt Ratio (%)	4.3	4.4	4.4	6	5.3	4.3
10	Mean ST Foreign Debt Ratio (%)	1.6	1.8	1.8	2.2	1.9	1.9
11	Mean Exports/Sales Ratio (%)	7.7	7.1	6	5.2	5.6	4.7
12	Fraction of Exporters (%)	20.1	17.3	16.3	14	13.6	13.2
13	Fraction of Firms with Foreign Debt (%)	43.9	39.4	38.7	38	34.5	31.8
				Publicly-Listed Firms			
1	Number of Firms	881	959	988	1,046	1,049	1,064
2	Mean Age	22	22	21	22	22	22
3	Mean Real Assets	274	294	327	385	338	395
4	Median Real Assets	49	47	45	47	37	40
5	Median Real Sales Growth Rate (%)	14.8	14.5	9.7	7	-7.8	16.7
6	Median Profit Rate (%)	3.9	3.6	3	2.4	2.8	7.1
7	Mean Leverage Ratio (%)	71	70.8	70.4	71.7	67.1	59
8	Mean ST Debt Ratio (%)	28.2	28.6	28.4	28.9	28.4	25.7
9	Mean Foreign Debt Ratio (%)	6.1	6.3	6.8	8.8	7.3	6.4
10	Mean ST Foreign Debt Ratio (%)	2.7	3	3.2	4	3.2	3.5
11	Mean Exports/Sales Ratio (%)	11.3	10.6	9.7	8.3	9.1	8.3
12	Fraction of Exporters (%)	30.5	28.5	26.7	22.9	22.3	19.5
13	Fraction of Firms with Foreign Debt (%)	66.5	64.7	64	62.2	60.5	53.9

Note: Real assets are in billion 1994 won. The profit rate is defined as the ratio of the pre-tax profit and the previous-year sales. The leverage ratio is defined as total liabilities over total assets. The ST debt ratio is defined as the amount of debt with original maturity less than or equal to one year divided by total liabilities. The foreign debt ratio is defined as foreign debt as a share of total liabilities. The foreign ST debt ratio is defined as short-term foreign debt over total liabilities. We remove the top and bottom 1% observations in terms of the sales growth rate and the profit rate.

Table 6: Summary Statistics for Liquidated Firms

	Panel A:					
		1995	1996	1997	1998	1999
1	Number of Exited Firms	71	100	218	206	57
	Exit Rate (%)	2.2	2.5	4.8	3.9	1.0
		Characteristics of Year Before Exit				
		1994	1995	1996	1997	1998
2	Mean Age	11	9	10	12	12
3	Mean Real Assets	17	18	20	25	15
4	Median Real Assets	13	13	13	12	11
5	Median Real Sales Growth Rate (%)	11.3	17.1	5.4	3.2	-41.6
6	Median Profit Rate (%)	0.4	-0.7	0.1	-4.4	-12.6
7	Mean Leverage Ratio (%)	92.7	98.3	96.3	104	112.7
8	Mean ST Debt Ratio (%)	41	38.2	38.5	41	36.2
9	Mean Foreign Debt Ratio (%)	2.5	1.3	3.3	3.3	3.9
10	Mean ST Foreign Debt Ratio (%)	1	0.6	1.1	1.2	0.8
11	Mean Exports/Sales Ratio (%)	6.6	3.4	3.5	3	2
12	Fraction of Exporters (%)	12.7	5	10.6	5.8	5.3
13	Fraction of Firms with Foreign Debt (%)	22.5	16	25.7	28.2	24.6
	Percentage of Exited Firms in					
	Industry 1	4	2	1	2	0
	Industry 2	34	41	35	32	32
	Industry 3	51	44	54	44	54
	Industry 4	7	8	8	15	5
	Industry 5	4	5	2	7	9
	Panel B:	**Extensive versus Intensive Margin**				
	Aggregate Sales Growth	18.06	12.49	10.19	-7.21	11.53
	Intensive Margin	18.39	12.89	11.22	-5.91	11.69
	Extensive Margin	-0.34	-0.4	-1.03	-1.3	-0.16
	% of aggregate sales growth	-1.87	-3.21	-10.21	18.01	-1.41

Note: Exited firms are firms that are liquidated. The exit rate of year t is computed as the ratio of the number of exited firms in year t and the sum of the number of firms survived from year t-1 to t and the number of firms exited in year t. Characteristic statistics of exited firms are reported for the year preceding the liquidation. The profit rate for exited firms is computed as the ratio of the after-tax profit and the previous-year sales since we don't have pre-tax profits for exited firms. The aggregate sales growth is computed as the total sales of all surviving firms in period t+1 divided by the total sales of both surviving and exited firms in period t. We measure the intensive margin with the ratio of the total sales of surviving firms in year t+1 and t, and measure the extensive margin as the ratio of the total year-t sales of exited firms and the total sales of all firms in year t. Industry 1 is Agriculture, Forestry, Fishing, and Mining; Industry 2 is Construction and Utility; Industry 3 is Manufacturing; Industry 4 is Wholesale and Retail Trade and Transportation; Industry 5 is Other Services.

Table 7: Cross-Section Regressions for Publicly-Listed Firms

Dependent Variable:	Crisis			Pre-Crisis		
Sales Growth	1	2	3	4	5	6
Chaebol Dummy	0.098**	0.101**	0.098**	0.086**	0.091**	0.088**
	(0.038)	(0.041)	(0.041)	(0.037)	(0.037)	(0.038)
Age	0.000	0.000	0.000	-0.002***	-0.002***	-0.002***
	(0.001)	(0.002)	(0.002)	(0.001)	(0.001)	(0.001)
Size	-0.077***	0.031	0.033	-0.026**	-0.017	-0.019
	(0.017)	(0.090)	(0.090)	(0.011)	(0.037)	(0.036)
Leverage Ratio	0.248**	4.15	4.265	-0.031	0.837	0.713
	(0.117)	(3.286)	(3.272)	(0.047)	(0.831)	(0.810)
Size * Leverage Ratio		-0.162	-0.167		-0.035	-0.030
		(0.132)	(0.132)		(0.034)	(0.033)
ST Debt Ratio	-0.204	-1.426	-1.434	-0.032	-1.705	-1.484
	(0.129)	(1.809)	(1.907)	(0.076)	(1.365)	(1.503)
Size * ST Debt Ratio		0.053	0.058		0.069	0.061
		(0.072)	(0.077)		(0.055)	(0.061)
Exports/Sales Ratio	0.211***	0.185***	0.163**	-0.038	-0.032	-0.035
	(0.060)	(0.060)	(0.064)	(0.047)	(0.049)	(0.049)
Foreign Debt Ratio	0.623**	-0.503		0.017	0.348	
	(0.302)	(4.539)		(0.163)	(2.430)	
Size * Foreign Debt Ratio		0.044			-0.014	
		(0.174)			(0.095)	
ST Foreign Debt Ratio			-4.464			-2.920
			(3.325)			(1.849)
Size * ST Foreign Debt Ratio			0.180			0.109
			(0.131)			(0.074)
LT Foreign Debt Ratio			1.982			1.316
			(9.503)			(2.917)
Size * LT Foreign Debt Ratio			-0.029			-0.047
			(0.361)			(0.116)
Observations	988	988	988	881	881	881
R-squared	0.158	0.176	0.189	0.156	0.160	0.164

Note: The dependent variable is firm sales growth between 1997 and 1998 for the crisis regressions and sales growth rate between 1995 and 1996 for the pre-crisis regressions. The independent variables are for year 1996 in the crisis regressions and for year 1994 in the pre-crisis regressions. Robust standard errors are reported in parentheses. *** denotes a p-value less than 1%, ** denotes a p-value less than 5%, and * denotes a p-value less than 10%. All regressions include industry dummies at the two-digit level.

Table 8: Cross-Section Regressions for the Full Sample

Dependent Variable: Sales Growth	Crisis 1	Crisis 2	Crisis 3	Pre-Crisis 4	Pre-Crisis 5	Pre-Crisis 6
Chaebol Dummy	0.095***	0.087***	0.087***	0.076***	0.076***	0.072***
	(0.027)	(0.027)	(0.027)	(0.022)	(0.022)	(0.024)
Age	0.000	0.000	0.000	-0.002***	-0.002***	-0.002***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Size	-0.055***	0.002	-0.001	-0.021***	0.003	0.003
	(0.007)	(0.025)	(0.025)	(0.006)	(0.023)	(0.022)
Leverage Ratio	0.105***	1.490*	1.517*	0.050	1.425**	1.363**
	(0.038)	(0.883)	(0.882)	(0.035)	(0.607)	(0.613)
Size * Leverage Ratio		-0.059	-0.060		-0.058**	-0.055**
		(0.037)	(0.037)		(0.025)	(0.026)
ST Debt Ratio	-0.026	1.380*	1.172	0.017	-1.445*	-1.214
	(0.044)	(0.801)	(0.820)	(0.059)	(0.842)	(0.867)
Size * ST Debt Ratio		-0.060*	-0.05		0.062*	0.052
		(0.034)	(0.035)		(0.035)	(0.036)
Exports/Sales Ratio	0.192***	0.182***	0.182***	-0.030	-0.028	-0.028
	(0.032)	(0.032)	(0.031)	(0.028)	(0.028)	(0.028)
Foreign Debt Ratio	0.417***	-2.817**		0.138	-1.297	
	(0.109)	(1.326)		(0.096)	(1.782)	
Size * Foreign Debt Ratio		0.134**			0.059	
		(0.053)			(0.074)	
ST Foreign Debt Ratio			-4.600***			-5.501
			(1.655)			(3.597)
Size * ST Foreign Debt Ratio			0.194***			0.228
			(0.067)			(0.154)
LT Foreign Debt Ratio			-2.705			0.990
			(2.173)			(2.238)
Size * LT Foreign Debt Ratio			0.141			-0.034
			(0.087)			(0.092)
Observations	4,285	4,285	4,285	3,151	3,151	3,151
R-squared	0.093	0.099	0.102	0.059	0.063	0.064

Note: The dependent variable is firm sales growth between 1997 and 1998 for the crisis regressions and sales growth rate between 1995 and 1996 for the pre-crisis regressions. The independent variables are for year 1996 in the crisis regressions and for year 1994 in the pre-crisis regressions. Robust standard errors are reported in parentheses. *** denotes a p-value less than 1%, ** denotes a p-value less than 5%, and * denotes a p-value less than 10%. All regressions include industry dummies at the two-digit level.

Table 9: Joint Distribution of Foreign Debt and Export Status

	Firms with Foreign Debt: Total Number	Firms with Foreign Debt: Fraction of Exporters	Firms with Foreign Debt: Fraction of Non-Exporters	Firms with No Foreign Debt: Total Number	Firms with No Foreign Debt: Fraction of Exporters	Firms with No Foreign Debt: Fraction of Non-Exporters
Full Sample	1,660	28.73	71.27	2,625	8.5	91.5
Asset bins						
1	83	10.84	89.16	988	5.06	94.94
2	384	24.22	75.78	687	11.64	88.36
3	497	30.38	69.62	574	9.93	90.07
4	696	32.18	67.82	376	9.57	90.43
Publicly-listed Firms	632	34.18	65.82	356	13.48	86.52

Note: A firm is classified as an exporter if its export-sales ratio is positive, and as a non-exporter otherwise. A firm is classified as a foreign debt holder if its foreign debt holdings are positive, and as a non-foreign-debt holder otherwise.

Table 10: Coefficients in Logit Exit Regressions

	Crisis			Pre-Crisis		
	1	2	3	4	5	6
Chaebol Dummy	-0.654*	-0.540	-0.531			
	(0.364)	(0.365)	(0.365)			
Age	-0.033***	-0.033***	-0.032***	-0.049***	-0.049***	-0.049***
	(0.008)	(0.008)	(0.008)	(0.011)	(0.011)	(0.011)
Size	0.086*	0.153***	0.149***	-0.173***	-0.114*	-0.103
	(0.046)	(0.048)	(0.048)	(0.065)	(0.067)	(0.067)
Leverage Ratio	2.278***	2.176***	2.168***	0.815***	0.787***	0.786***
	(0.397)	(0.389)	(0.387)	(0.302)	(0.300)	(0.302)
ST Debt Ratio	1.691***	1.678***	1.727***	0.921**	0.909**	0.875**
	(0.289)	(0.291)	(0.296)	(0.420)	(0.420)	(0.425)
Profit/Assets	-2.066*	-2.160*	-2.145*	-3.637***	-3.658***	-3.681***
	(1.234)	(1.186)	(1.190)	(0.935)	(0.936)	(0.936)
Exports/Sales Ratio	-0.576	-0.509	-0.509	0.364	0.456	0.455
	(0.460)	(0.455)	(0.457)	(0.549)	(0.541)	(0.543)
Foreign Debt Ratio	-0.038	61.090***		-0.874	60.510***	
	(0.725)	(12.880)		(1.866)	(22.410)	
Size * Foreign Debt Ratio		-2.608***			-2.669***	
		(0.555)			(0.994)	
ST Foreign Debt Ratio			72.640***			92.060*
			(27.160)			(52.370)
Size * ST Foreign Debt Ratio			-3.146***			-3.944*
			(1.173)			(2.254)
LT Foreign Debt Ratio			48.980***			61.000**
			(16.990)			(30.740)
Size * LT Foreign Debt Ratio			-2.058***			-2.731*
			(0.742)			(1.412)
Observations	4,696	4,696	4,696	3,398	3,398	3,398
Pseudo R-squared	0.156	0.162	0.162	0.135	0.139	0.139

Note: The dependent variable is either 1 if the firm exits in 1997 and 1998 or 0 otherwise for the crisis regressions, and is either 1 if the firm exits in 1995 and 1996 or 0 otherwise in the pre-crisis regressions. The independent variables are for 1996 in the crisis regressions and for 1994 for the pre-crisis regressions. Robust standard errors are reported in parentheses. *** denotes a p-value less than 1%, ** denotes a p-value less than 5%, and * denotes a p-value less than 10%. All regressions include industry dummies at the one-digit level.

Table 11: Differential Impact of Short-term Foreign Debt

	Intensive margin: sales growth			Extensive margin: exit probability			Both margins: sales growth		
Asset bins	Observed	Counterfactual foreign debt	Counterfactual export sales	Observed	Counterfactual foreign debt	Counterfactual export sales	Observed	Counterfactual foreign debt	Counterfactual export sales
	1	2	3	4	5	6	7	8	9
1	-2.9%	-4.5%	-3.2%	7.2%	14.6%	7.3%	-10.1%	-17.7%	-10.5%
2	-4.7%	-5.4%	-6.0%	7.6%	10.5%	7.8%	-12.8%	-15.4%	-14.1%
3	-8.5%	-7.8%	-9.9%	6.9%	7.7%	7.1%	-15.0%	-15.0%	-16.4%
4	-13.5%	-9.0%	-15.7%	6.2%	4.5%	6.4%	-16.2%	-10.5%	-18.5%
Total	-12.6%	-8.6%	-14.6%	7.0%	9.3%	7.1%	-15.9%	-11.1%	-17.9%

Note: Firms are classified into four equal number bins according to their size. The smallest firms are in asset bin 1, and the largest firms are in bin 4. The statistics under the intensive margin are predicted sales growth using the regression estimates in column 3 of Table 8. The statistics under the extensive margin are predicted exit probabilities computed using column 3 of Table 10. The sales growth under both margins is computed as the average of the predicted sales growth rate conditional on survival and the sales growth rate of -1 conditional on exit weighted by the survival and exit probability, respectively. The columns under observed characteristics are computed using the 1996 observables for each firm. The columns under counterfactual foreign debt are computed for the scenario in which each firm has a 10 percent hypothetical increase of its leverage ratio through increasing short-term foreign debt holdings and all the other variables as observed in the data. The columns under counterfactual export sales are computed for the scenario in which all firms have zero export-sales ratios and all the other variables as observed in the data.

Table 12: Profit Regressions for Publicly Listed Firms

Dependent Variable:	Crisis			Pre-Crisis		
Profit/$Sales_{-1}$	1	2	3	4	5	6
Chaebol Dummy	0.024	0.024	0.023	0.001	0.004	0.003
	(0.024)	(0.025)	(0.025)	(0.011)	(0.011)	(0.011)
Age	0.001	0.001	0.001	-0.001**	-0.001**	-0.001**
	(0.001)	(0.001)	(0.001)	(0.000)	(0.000)	(0.000)
Size	-0.036***	-0.008	-0.007	-0.009**	0.018	0.019
	(0.006)	(0.019)	(0.020)	(0.004)	(0.017)	(0.017)
Leverage	-0.091***	0.687	0.708	-0.040*	0.895*	0.889*
	(0.034)	(0.570)	(0.570)	(0.024)	(0.517)	(0.523)
Size * Leverage		-0.032	-0.033		-0.038*	-0.038*
		(0.024)	(0.024)		(0.021)	(0.021)
ST Debt Ratio	-0.063**	0.496	0.494	-0.071***	0.018	0.096
	(0.031)	(0.436)	(0.473)	(0.023)	(0.302)	(0.335)
Size * ST Debt Ratio		-0.023	-0.022		-0.004	-0.007
		(0.018)	(0.020)		(0.012)	(0.014)
Exports/Sales	0.133***	0.125***	0.121***	0.025*	0.025*	0.025*
	(0.021)	(0.021)	(0.021)	(0.014)	(0.014)	(0.014)
Foreign Debt Ratio	0.205***	-0.748		-0.013	-0.388	
	(0.055)	(0.753)		(0.046)	(0.480)	
Size * Foreign Debt Ratio		0.038			0.015	
		(0.030)			(0.019)	
ST Foreign Debt Ratio			-1.443			-0.369
			(0.962)			(0.716)
Size * ST Foreign Debt Ratio			0.062			0.017
			(0.039)			(0.029)
LT Foreign Debt Ratio			-0.314			-0.179
			(1.438)			(0.760)
Size * LT Foreign Debt Ratio			0.026			0.004
			(0.058)			(0.030)
Observations	988	988	988	881	881	881
R-squared	0.182	0.189	0.192	0.19	0.202	0.204

Note: The dependent variable is firm profits in 1998 as a share of sales in 1997 for the crisis regressions and firm profits in 1996 as a share of sales in 1995 for the pre-crisis regressions. The independent variables are for year 1996 in the crisis regressions and for year 1994 in the pre-crisis regressions. Robust standard errors are reported in parentheses. *** denotes a p-value less than 1%, ** denotes a p-value less than 5%, and * denotes a p-value less than 10%. All regressions include industry dummies at the two-digit level.

Table 13: Profit Regressions for the Full Sample

Dependent Variable:	Crisis			Pre-Crisis		
$Profit/Sales_{-1}$	1	2	3	1	2	3
Chaebol Dummy	-0.015	-0.013	-0.014	0.015*	0.017*	0.016*
	(0.018)	(0.018)	(0.018)	(0.009)	(0.009)	(0.009)
Age	0.000	0.000	0.000	0.000	0.000	0.000
	0.000	0.000	0.000	0.000	0.000	0.000
Size	-0.022***	0.023**	0.0242**	-0.008***	0.017	0.018
	(0.003)	(0.010)	(0.010)	(0.003)	(0.016)	(0.016)
Leverage	-0.129***	1.012***	1.013***	-0.122***	0.457	0.454
	(0.018)	(0.292)	(0.292)	(0.025)	(0.483)	(0.479)
Size * Leverage		-0.049***	-0.049***		-0.024	-0.024
		(0.013)	(0.013)		(0.021)	(0.021)
ST Debt Ratio	-0.054***	0.702***	0.745***	-0.054***	0.521*	0.576*
	(0.014)	(0.209)	(0.217)	(0.018)	(0.297)	(0.304)
Size * ST Debt Ratio		-0.032***	-0.034***		-0.024*	-0.027**
		(0.009)	(0.010)		(0.013)	(0.013)
Exports/ Sales	0.082***	0.075***	0.075***	-0.007	-0.009	-0.008
	(0.011)	(0.011)	(0.011)	(0.010)	(0.010)	(0.010)
Foreign Debt Ratio	0.090***	-0.424		-0.033	0.018	
	(0.030)	(0.412)		(0.037)	(0.453)	
Size * Foreign Debt Ratio		0.021			-0.002	
		(0.017)			(0.018)	
ST Foreign Debt Ratio			-1.032*			0.263
			(0.571)			(0.547)
Size * ST Foreign Debt Ratio			0.0465*			-0.009
			(0.024)			(0.023)
LT Foreign Debt Ratio			0.159			0.141
			(0.660)			(0.674)
Size * LT Foreign Debt Ratio			-0.004			-0.01
			(0.028)			(0.028)
Observations	4,285	4,285	4,285	3,151	3,151	3,151
R-squared	0.132	0.145	0.145	0.120	0.125	0.126

Note: The dependent variable is firm profits in 1998 as a share of sales in 1997 for the crisis regressions and firm profits in 1996 as a share of sales in 1995 for the pre-crisis regressions. The independent variables are for year 1996 in the crisis regressions and for year 1994 in the pre-crisis regressions. Robust standard errors are reported in parentheses. *** denotes a p-value less than 1%, ** denotes a p-value less than 5%, and * denotes a p-value less than 10%. All regressions include industry dummies at the two-digit level.

Table 14: Investment Regressions for Publicly-listed Firms

Dependent Variable:	Crisis			Pre-Crisis		
I/K_{-1}	1	2	3	1	2	3
Chaebol Dummy	0.051	0.052	0.052	0.190***	0.193***	0.190***
	(0.045)	(0.045)	(0.045)	(0.068)	(0.069)	(0.070)
Age	-0.001	-0.001	-0.001	-0.003***	-0.003***	-0.003***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Size	-0.008	0.079	0.08	-0.008	0.027	0.027
	(0.011)	(0.051)	(0.051)	(0.015)	(0.035)	(0.035)
Leverage	0.007	2.947*	2.958*	-0.050	1.528*	1.542*
	(0.086)	(1.649)	(1.655)	(0.035)	(0.893)	(0.906)
Size * Leverage		-0.121*	-0.122*		-0.065*	-0.065*
		(0.067)	(0.067)		(0.036)	(0.037)
ST Debt Ratio	0.023	-0.414	-0.388	0.037	-0.274	-0.27
	(0.075)	(1.192)	(1.282)	(0.085)	(1.391)	(1.532)
Size * ST Debt Ratio		0.019	0.018		0.012	0.011
		(0.049)	(0.053)		(0.057)	(0.063)
Exports/ Sales	0.076	0.057	0.056	-0.077	-0.075	-0.068
	(0.058)	(0.045)	(0.046)	(0.066)	(0.067)	(0.067)
Foreign Debt Ratio	-0.146	-0.760		0.083	-2.355	
	(0.096)	(1.543)		(0.173)	(1.928)	
Size * Foreign Debt Ratio		0.024			0.098	
		(0.062)			(0.078)	
ST Foreign Debt Ratio			-1.255			-0.457
			(2.030)			(3.250)
Size * ST Foreign Debt Ratio			0.042			0.032
			(0.083)			(0.130)
LT Foreign Debt Ratio			-0.371			-3.777
			(2.912)			(2.884)
Size * LT Foreign Debt Ratio			0.010			0.144
			(0.118)			(0.117)
Lagged I/K_{-1}	0.081***	0.076**	0.076**	0.175***	0.167***	0.167***
	(0.031)	(0.031)	(0.031)	(0.047)	(0.047)	(0.048)
Observations	949	949	949	864	864	864
R-squared	0.067	0.085	0.086	0.185	0.189	0.191

Note: The dependent variable is firm real investment in 1998 as a share of the replacement value of real capital stock in 1997 for the crisis regressions and firm real investment in 1996 as a share of the replacement value of real capital stock in 1995 for the pre-crisis regressions. The independent variables are for year 1996 in the crisis regressions and for year 1994 in the pre-crisis regressions. Robust standard errors are reported in parentheses. *** denotes a p-value less than 1%, ** denotes a p-value less than 5%, and * denotes a p-value less than 10%. All regressions include industry dummies at the two-digit level.

Table 15: Investment Regressions for the Full Sample

Dependent Variable:	Crisis			Pre-Crisis		
I/K_{-1}	1	2	3	1	2	3
Chaebol Dummy	0.056*	0.048	0.049	0.182***	0.175***	0.174***
	(0.030)	(0.031)	(0.031)	(0.038)	(0.038)	(0.038)
Age	0.000	0.000	0.000	-0.002***	-0.002***	-0.002***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Size	-0.015**	-0.013	-0.013	-0.023***	-0.007	-0.006
	(0.006)	(0.023)	(0.024)	(0.008)	(0.023)	(0.023)
Leverage	0.018	0.567	0.564	-0.034	1.344**	1.345**
	(0.029)	(0.710)	(0.710)	(0.035)	(0.628)	(0.633)
Size * Leverage		-0.023	-0.023		-0.058**	-0.058**
		(0.030)	(0.030)		(0.026)	(0.026)
ST Debt Ratio	-0.018	-1.041	-1.048	-0.038	-1.662*	-1.593
	(0.035)	(0.705)	(0.730)	(0.044)	(0.959)	(1.010)
Size * ST Debt Ratio		0.044	0.044		0.0689*	0.065
		(0.030)	(0.031)		(0.041)	(0.043)
Exports/ Sales	0.072*	0.070*	0.070*	-0.009	-0.006	-0.004
	(0.039)	(0.038)	(0.038)	(0.035)	(0.035)	(0.035)
Foreign Debt Ratio	0.032	-2.185*		-0.067	-4.308***	
	(0.075)	(1.287)		(0.091)	(1.330)	
Size * Foreign Debt Ratio		0.092*			0.175***	
		(0.054)			(0.055)	
ST Foreign Debt Ratio			-1.552			-2.763
			(2.341)			(2.789)
Size * ST Foreign Debt Ratio			0.067			0.122
			(0.096)			(0.115)
LT Foreign Debt Ratio			-2.633			-4.903**
			(1.868)			(1.977)
Size * LT Foreign Debt Ratio			0.109			0.193**
			(0.079)			(0.082)
Lagged I/K_{-1}	0.088***	0.088***	0.088***	0.164***	0.161***	0.161***
	(0.019)	(0.019)	(0.019)	(0.027)	(0.027)	(0.027)
Observations	4,072	4,072	4,072	2,970	2,970	2,970
R-squared	0.043	0.046	0.046	0.097	0.102	0.103

Note: The dependent variable is firm real investment in 1998 as a share of the replacement value of real capital stock in 1997 for the crisis regressions and firm real investment in 1996 as a share of the replacement value of real capital stock in 1995 for the pre-crisis regressions. The independent variables are for year 1996 in the crisis regressions and for year 1994 in the pre-crisis regressions. Robust standard errors are reported in parentheses. *** denotes a p-value less than 1%, ** denotes a p-value less than 5%, and * denotes a p-value less than 10%. All regressions include industry dummies at the two-digit level.

CHAPTER III

How Reliable Are Local Projection Estimators of Impulse Responses?

with Lutz Kilian

3.1 Introduction

Estimates of structural impulse response functions are of central interest in empirical macroeconomics. Conventional estimates of structural impulse responses are based on vector autoregressive (VAR) models. VAR impulse response coefficients are nonlinear functions of the estimates of the VAR slope parameters and of the VAR innovation variance-covariance matrix. It is well known that the finite-sample accuracy of conventional asymptotic and bootstrap approximations to the distribution of the impulse response estimator is undermined by the bias of the impulse response estimator. This bias arises from two distinct sources: the small-sample bias of the estimates of the VAR slope parameters and the additional bias induced by the non-linear transformations of the estimated parameters.

Over the last two decades, several econometric methods have been proposed to address this problem. For example, the bias-corrected bootstrap confidence interval of Kilian (1998a) explicitly accounts for the first source of bias by constructing bias-corrected slope coefficient estimates. This bias-adjusted bootstrap provides a significant improvement in coverage accuracy over the standard bootstrap of Runkle (1987) and the asymptotic delta method interval of Lütkepohl (1990), but it does not necessarily yield accurate coverage when the process is highly persistent or when the model includes deterministic time trends.[48] A large number of studies has investigated the small-sample and asymptotic accuracy of these methods (see, e.g., Griffiths and Lütkepohl (1993), Fachin and Bravetti (1996). Kilian (1998a,b,c; 1999), Berkowitz and Kilian (2000), Benkwitz, Lütkepohl and Neumann (2000), and Kilian (2001)).[49]

Kilian and Chang (2000) demonstrated that the coverage accuracy of all traditional asymptotic and bootstrap methods, including bias-adjusted methods, tends to deteriorate in large-dimensional VAR models at longer horizons, when the data are highly persistent. This finding motivated the use of nonstandard asymptotic approximations based on local-to-unity models. Since the bias of the impulse responses typically worsens, as the dominant autoregressive root approaches unity, that approach seemed well suited to dealing with the bias problem. For example, Wright (2000), building on Stock (1991), proposed conservative asymptotic impulse response confidence intervals based on local-to-unity approximations to the largest root of the autoregressive process. His method proved computationally intractable in VAR models, however. In related work, Gospodinov (2004) proposed an asymptotic impulse response confidence interval based on the inversion of the likelihood ratio statistic. His method differs from Wright (2000) in that his interval is not conservative, but asymptotically exact. Gospodinov's approach, however, is limited to univariate autoregressions and requires knowledge of a point null which is unlikely to be available in practice. As yet another alternative, Hansen (1999) proposed a grid bootstrap method for the dominant root in univariate autoregressive processes. This grid bootstrap provides correct asymptotic coverage regardless of whether the autoregressive model is near-integrated or exactly integrated. Gospodinov (2004), however, reports that the coverage accuracy of the grid bootstrap method applied to impulse responses is too low at short horizons. Pesavento and Rossi (2006) were the first to propose

[48]An alternative bias-adjusted method for impulse responses in univariate autoregressions has been proposed by Andrews and Chen (1994). Related work also includes Rudebusch (1992).

[49]Related studies focusing on univariate autoregressions include Berkowitz, Birgean and Kilian (1999), Inoue and Kilian (2002a, 2003), and Pesavento and Rossi (2007).

operational confidence intervals for VAR impulse responses based on a local-to-unity approximation to the asymptotic distribution of that estimator. Their method is designed to achieve accurate coverage at long horizons when the data are highly persistent. At short horizons, its coverage accuracy may not be satisfactory. Pesavento and Rossi (2006) also proposed another method designed to yield accurate coverage at both short and long horizons, but that method may be conservative at medium horizons. Thus, even twenty years after the first confidence intervals were developed for VAR impulse responses, there is no single method that resolves the bias problem in all situations.

Our paper explores a new idea for dealing with this bias problem that is motivated by recent advances in estimating impulse response functions from local projections. Local (linear) projections (LPs) were proposed by Jordà (2005, 2007) on the grounds that such projections may be more robust to model misspecification. We instead focus on another potentially attractive feature of that method, namely its ability to ameliorate the bias problem that has undermined the accuracy of inference on impulse responses in practice. The basic idea of the LP method is that we directly estimate a sequence of linear projections of the future value of the dependent variable on the current information set. This approach will be asymptotically equivalent to the VAR-based approach, provided the data generating process is stationary and linear. Unlike VAR impulse response estimates, however, impulse responses based on local projections do not require any nonlinear transformations of the estimated slope parameters. Rather the slope parameters themselves are the reduced-form impulse response coefficients and - with the help of any estimate of the structural impact multiplier matrix - can easily be used to construct the structural impulse response function. Since local projections do not involve any nonlinear transformation, they are likely to be better approximated by Gaussian distributions. A very similar point has been made by Davidson (2000) in the context of bootstrap methods. Thus, local projection methods have the potential of greatly reducing the bias of impulse response point estimates and of increasing the coverage accuracy of impulse response confidence intervals.

The LP method is not without drawbacks, however. One of its potential disadvantages is that the LP estimator tends to have higher variance, when the data generating process is well approximated by a vector autoregression, since local projections impose less structure on the estimation problem. This increase in the variance would be expected to increase the mean-squared error (MSE) of impulse response point estimates and the average length of LP impulse response confidence intervals. Moreover, nothing is known about the extent of small-sample bias in the slope parameters of local projections compared with the well-known small-sample bias in the VAR slope parameters (see Pope 1990). Hence, the question of whether local projections should replace VAR models in constructing structural impulse response point and interval estimates is ultimately an empirical question.

In this paper, we address this question by comparing the finite-sample properties of impulse response confidence intervals estimated by VAR models and by local projections using simulation studies. Our objective is to provide some practical guidance about which of the combinations of estimation method (VAR versus LP) and method of inference (asymptotic versus bootstrap) is likely to be most reliable in practice, when the VAR framework is considered a good approximation to the data generating process. Throughout the paper we will maintain the assumption of stationarity. A summary of the alternative approaches considered in this paper is provided in Table 16.

We make three distinct contributions. First, for each method, we compare the pointwise coverage accuracy and average length of the asymptotic and bootstrap confidence interval in stationary models. Our analysis allows for model specification uncertainty. Specifically, for the VAR model we focus on (1) the asymptotic delta method interval of Lütkepohl (1990) and (2) percentile intervals based on the bias-corrected bootstrap method proposed by Kilian (1998a,b; 1999). For local projections, in contrast, we investigate (3) the asymptotic interval proposed by Jordà (2005, 2007) and (4) we propose a bootstrap percentile interval based on a suitably designed block bootstrap method.[50] We provide some intuition for the relative performance of these approaches by investigating the bias, standard deviation, and MSE of the impulse response point estimators. Our simulation analysis is based on a commonly used stylized bivariate VAR(1) data generating process as well as an additional high-dimensional VAR data generating process of the type used in studying responses to monetary policy shocks. Such realistic models are rarely analyzed in the literature, given the computational costs of simulating large-dimensional VAR systems.[51]

We find that, contrary to our conjecture, the LP point estimator tends to be both more biased and more variable than VAR based estimator, resulting in often excessively wide LP intervals with less than nominal coverage. Given its greater average width, the asymptotic LP interval tends to be more accurate than the VAR delta method interval. It also is more accurate than the bootstrap percentile LP interval, although again at the cost of added width. The reason is that standard applications of the bootstrap tend to reinforce the small-sample bias, much like in the Runkle (1987) method (see Kilian 1998a). Unlike in the VAR context, there is no obvious way of correcting for this bias in the LP estimator. The bias-adjusted VAR bootstrap interval, in contrast, tends to be more accurate than either LP interval in small samples and typically shorter on average. It also is more accurate than traditional asymptotic delta method intervals and only slightly wider on average, consistent with previous findings in the literature. Thus, for pointwise impulse response intervals, among the

[50] Although Jordà (2007) discusses the potential benefits from bootstrapping local projections, he does not explore bootstrap methods in his work. Our bootstrap percentile interval provides an alternative to asymptotic inference. We do not investigate the use of percentile-t intervals for the LP method since it is not clear how to estimate the variance of the bootstrap impulse response estimator.

[51] Kilian and Chang (2000) is a notable exception.

methods considered, no method is more accurate than the bias-adjusted bootstrap method for VAR models. For larger sample sizes, the LP asymptotic interval and the VAR bootstrap interval have similar coverage accuracy, but the LP interval tends to be much wider on average. Additional simulation evidence suggests that similar results hold even for VARMA data generating processes, when the finite lag order VAR model is only an approximation.

Second, in addition to evaluating the pointwise coverage accuracy of the intervals at each impulse response horizon, we also evaluate the coverage accuracy of the joint confidence interval proposed by Jordà (2007) whose small-sample properties have not been examined previously. These joint intervals can be computed from either local projections or VAR models. Joint confidence intervals are of particular interest to empirical researchers who want to evaluate the uncertainty about the entire path of the impulse response function. We focus on the asymptotic interval estimates (5) and (6). Our simulation results suggest that the coverage accuracy of the joint intervals can be erratic, regardless of the method used. Thus, joint intervals have to be used with caution. Our analysis also shows that neither interval is uniformly preferred over the other, although the joint LP interval is invariably much wider on average. Bootstrapping the standard error does not improve the accuracy of the asymptotic LP interval.

Third, we illustrate the practical differences that may arise from the choice of different methods of estimation and inference by re-examining the empirical findings for a standard monthly VAR model of monetary policy. Using data for 1970.1-2007.12, the VAR model provides evidence that monetary policy contractions cause a temporary reduction in output and a temporary drop in real commodity prices, but we find no evidence that inflation is significantly reduced. Despite the inclusion of commodity prices, the estimates exhibit the well known price puzzle. The use of joint VAR intervals does not change any of these conclusions, although it affects the degree of statistical significance especially at longer horizons. Since most statistically significant impulse response estimates are obtained at shorter horizons, when pointwise and joint intervals tend to be similar, we conclude that typically the use of joint intervals will not overturn the substantive findings of VAR studies based on pointwise intervals. Likewise, estimates from the LP and VAR method are qualitatively similar, but the LP estimates tend to be more erratic and less precisely estimated. This tendency is even more pronounced if we focus on subsamples such as the post-Volcker period. Our analysis provides no compelling reason to abandon traditional VAR methods of constructing impulse response estimates in favor of the LP method.

The remainder of the paper is organized as follows. Section 3.2 briefly establishes the notation and contrasts the construction of impulse response estimates from VAR models and from local projections. In section 3.3, we build intuition based on results from a Monte Carlo study that employs a stylized bivariate VAR(1) data-generating process used in the previous literature. The simulation results for a more realistic VAR(12) model are presented in section 3.4 along with a comparison of the empirical estimates. Section 3.5 contains some preliminary simulation results for infinite-order VAR processes. We conclude in Section 3.6.

3.2 Review of VARs and Local Projections

3.2.1 Data-Generating Process

Consider a K-dimensional linear vector autoregressive data-generating processes (DGP) of finite order p:[52]

$$y_t = B_1 y_{t-1} + \cdots + B_p y_{t-p} + e_t, \tag{16}$$

[52]This approach is standard in the literature. Alternatively, one could interpret a vector autoregression as an approximation to general stationary linear DGP (see, e.g., Lütkepohl and Poskitt 1991; Inoue and Kilian 2002). This case will be addressed in section 3.5.

where $t = p+1, ..., T$, $y_t = (y_{1t}, \cdots, y_{Kt})'$ is a $(K \times 1)$ random vector, $B_i, i = 1, ..., p$, are $(K \times K)$ coefficient matrices and $e_t = (e_{1t}, \cdots, e_{Kt})'$ is K-dimensional i.i.d. white noise, i.e., $E(e_t) = 0$, $E(e_t e_s') = 0$ for $s \neq t$ and $E(e_t e_t') = \Sigma_e$ where Σ_e is non-singular and positive definite.[53] All values of z satisfying $\det(I_K - B_1 z - \cdots - B_p z^p) = 0$ lie outside the unit circle. For expository purposes, we abstract from deterministic regressors, although we will allow for an intercept in estimation throughout this paper. This VAR process can be written in structural form as:

$$A_0 y_t = A_1 y_{t-1} + \cdots + A_p y_{t-p} + \varepsilon_t, \tag{17}$$

where $\Sigma_\varepsilon = I_K$ without loss of generality.

3.2.2 Impulse Responses

Impulse responses to VAR reduced-form disturbances are obtained recursively as

$$\Phi_h^{VAR(p)} = \sum_{l=1}^{h} \Phi_{h-l}^{VAR(p)} B_l, \quad h = 1, 2, \ldots, H, \tag{18}$$

where $\Phi_0^{VAR(p)} = I_K$ and $B_l = 0$ for $l > p$. The corresponding responses to structural shocks are given by:

$$\Theta_h^{VAR(p)} = \Phi_h^{VAR(p)} A_0^{-1}, \quad h = 0, 1, \ldots, H, \tag{19}$$

where A_0^{-1} satisfies $A_0^{-1}(A_0^{-1})' = \Sigma_e$. For the purpose of the analysis below, we postulate that A_0^{-1} is a lower triangular matrix. Element (i, j) of $\Theta_h^{VAR(p)}$ is $\theta_{ij,h}^{VAR(p)}$ and represents the response of variable i to a one-time structural shock j, h periods ago. By construction, $\theta_{ij,h}^{VAR(p)}$ is a nonlinear function of B and Σ_e. Estimates $\hat{\Theta}_h^{VAR(p)}$ are

[53] The assumption of i.i.d. innovations is common in applied work and provides a useful benchmark for our purposes. It could be relaxed with suitable changes in the theory and implementation of the asymptotic and bootstrap approach (see Goncalves and Kilian 2004, 2007).

constructed by substituting the least-squares estimates of B and Σ_e obtained from regression (16).

An alternative approach to estimating reduced form impulse responses is to fit the linear projection

$$y_{t+h} = \mu + F_1 y_t + F_2 y_{t-1} + \cdots + F_q y_{t-q+1} + u_{t+h} \quad \text{for } h = 1, \ldots, H, \tag{20}$$

where u_t may be serially correlated or heteroskedastic (see Jordà 2005, 2007). The lag length q needs not be common across different horizons. By construction, the slope F_1 can be interpreted as the response of y_{t+h} to a reduced-form disturbance in period t:

$$\Phi_h^{LP(q)} = F_1 = E(y_{t+h}|e_t = 1; y_t, \ldots, y_{t-q}) - E(y_{t+h}|e_t = 0; y_t, \ldots, y_{t-q}), \quad h = 1, \ldots, H. \tag{21}$$

$\Phi_0^{LP(q)} = I_K$. The corresponding structural impulse responses are

$$\Theta_h^{LP(q)} = \Phi_h^{LP(q)} A_0^{-1}, \quad h = 0, 1, \ldots, H, \tag{22}$$

where A_0^{-1} is obtained based on the VAR model as described earlier.[54] $\theta_{ij,h}^{LP(q)}$ denotes the response of variable i to a one-time structural shock j, h periods ago. Estimates $\hat{\Theta}_h^{LP(q)}$ are constructed from the the VAR(p) estimate $\hat{A}_0^{-1}$ and the $\widehat{\Phi}_h^{LP(q)}$estimates obtained from a sequence of least-squares regressions (20) for each horizon h. Under the maintained assumption of the DGP in equation (16), both $\widehat{\theta}_{ij,h}^{VAR(p)}$ and $\widehat{\theta}_{ij,h}^{LP(q)}$ will be consistent for $\theta_{ij,h}^{VAR(p)}$.

[54] Jordà (2005) does not explicitly discuss the distinction between the structural and reduced-form impulse responses. The Gauss code provided by Jordà, however, shows that his structural impulse responses are constructed using the VAR estimate of A_0^{-1}.

3.2.3 Confidence Intervals

Asymptotic Confidence Intervals

Let $\beta = vec(B_1, B_2, \cdots, B_p)$ and $\sigma = vech(\Sigma_e)$. Under suitable moment restrictions, the asymptotic distribution of the VAR impulse response estimator can be derived by the delta method:

$$\sqrt{T} vec\left(\hat{\Theta}_h^{VAR(p)} - \Theta_h^{VAR(p)}\right) \xrightarrow{d} N\left(0, C_h \Sigma_{\hat{\beta}} C_h' + \overline{C}_h \Sigma_{\hat{\sigma}} \overline{C}_h'\right) \tag{23}$$

where $C_0 = 0$, $C_h = (A_0^{-1} \otimes I_K) G_h$ with $G_h = \partial vec(\Phi_h^{VAR(p)})/\partial\beta'$, and $\bar{C}_h = (I_K \otimes \Phi_h^{VAR(p)}) \partial vec(A_0^{-1})/\partial\sigma'$. Explicit expressions for the asymptotic variance of the impulse response estimator can be found in Lütkepohl (1990). The nominal $(1-\alpha)\%$ confidence interval satisfies

$$P\left(\hat{\theta}_{ij,h}^{VAR(p)} - z_{1-\alpha/2} \frac{1}{\sqrt{T}} \widehat{\sigma}\left(\widehat{\theta}_{ij,h}^{VAR(p)}\right) \leq \theta_{ij,h}^{VAR(p)} \leq \hat{\theta}_{ij,h}^{VAR(p)} + z_{1-\alpha/2} \frac{1}{\sqrt{T}} \widehat{\sigma}\left(\widehat{\theta}_{ij,h}^{VAR(p)}\right)\right) = 1-\alpha, \tag{24}$$

where $\widehat{\sigma}\left(\widehat{\theta}_{ij,h}^{VAR(p)}\right)$ is the square root of element $(K(j-1)+i, K(j-1)+i)$ of $\left(\widehat{C}_h \hat{\Sigma}_{\hat{\beta}} \widehat{C}_h' + \widehat{\overline{C}}_h \hat{\Sigma}_{\hat{\sigma}} \widehat{\overline{C}}_h'\right)$, $z_{1-\alpha/2}$ denotes the $(1-\alpha/2)$-quantile of the $N(0,1)$ distribution.

The asymptotic confidence interval of the corresponding LP estimator proposed by Jordà (2005) is

$$P\left(\hat{\theta}_{ij,h}^{LP(q)} - z_{1-\alpha/2} \frac{1}{\sqrt{T}} \widehat{\sigma}\left(\widehat{\theta}_{ij,h}^{LP(q)}\right) \leq \theta_{ij,h}^{LP(q)} \leq \hat{\theta}_{ij,h}^{LP(q)} + z_{1-\alpha/2} \frac{1}{\sqrt{T}} \widehat{\sigma}\left(\widehat{\theta}_{ij,h}^{LP(q)}\right)\right) = 1-\alpha. \tag{25}$$

Here $\widehat{\sigma}\left(\widehat{\theta}_{ij,h}^{LP(q)}\right)$ is the square root of element $(K(j-1)+i, K(j-1)+i)$ of

$$\left(A_0^{-1\prime} \otimes I_K\right)\left(\left(y_t' M_x y_t\right)^{-1} \otimes \widehat{\Sigma}_u\right)\left(A_0^{-1} \otimes I_K\right) + \overline{G}_h \Sigma_{\hat{\sigma}} \overline{G}_h', \tag{26}$$

where $M_x = I - X(X'X)^{-1}X'$, $X = \begin{bmatrix} 1 & y_{t-1} & y_{t-2} & \ldots & y_{t-q} \end{bmatrix}$, $\bar{G}_h = (I_K \otimes \Phi_h^{LP(q)}) \partial vec(A_0^{-1})/\partial\sigma'$, and $\widehat{\Sigma}_u = E(u_{t+h}u'_{t+h})$ in equation (20). The first additive component of this variance-covariance matrix captures the variance of $vec(\widehat{\Phi}_h^{LP(q)} A_0^{-1})$ and reflects the uncertainty associated with the slope parameter estimates. The second additive component incorporates the estimation uncertainty associated with the estimate of A_0^{-1}(see Jordà 2007).[55] Following Jordà (2005), we employ the Newey-West estimator of $\hat{\Sigma}_u$.[56]

Bootstrap Confidence Intervals

Confidence intervals can also be obtained by bootstrap approximations. For the VAR impulse response estimator, we consider the well-established bias-corrected bootstrap confidence interval proposed by Kilian (1998a, 1999). The reader is referred to the relevant literature for details of that procedure.[57] For the LP impulse response estimator, no bootstrap methods have been considered to date. Although Jordà (2007) discusses the potential benefits from bootstrapping the LP estimator, he does not explore any bootstrap methods in his work. In this paper, we propose a block bootstrap approach since the error term in LP regressions is serially correlated. By construction, the LP impulse response estimate for horizon h depends on the $(1+q)$ tuple $(y_{t+h}, y_t, y_{t-1}, \cdots, y_{t-q+1})$. To preserve the correlation in the data, we first construct the set of all possible $(1+q)$ tuples. Then blocks of l consecutive $(1+q)$ tuples

[55] Jordà (2005) abstracts from this second component. This causes the asymptotic interval for LP too narrow at short horizons. Additional simulation results show that adding the second term significantly improves coverage accuracy of the asymptotic confidence interval at short horizons, with a marginal increase in average length.

[56] Jordà (2005) shows that the disturbance terms in a local projection has a moving average component of order h under our assumptions: $u_{t+h} = e_{t+h} + \Phi_1^{VAR} e_{t+h-1} + \Phi_2^{VAR} e_{t+h-2} + \cdots + \Phi_{h-1}^{VAR} e_{t+1}$. This suggests that we set the truncation lag for the Newey-West estimator to be h for each local projection horizon h. While the results are not overly sensitive to the choice of the truncation lag, estimating Σ_u by least squares would seriously undermine the accuracy of the LP interval.

[57] We implement this method as discussed in Kilian (1998b,c, 1999) using the full double loop rather than using the computational short-cut proposed in Kilian (1998a). The first-order bias is estimated using the asymptotic closed-form solutions proposed by Pope (1990) rather than the bootstrap method. For a detailed description see, e.g., Kilian (1998b).

are drawn (see, e.g., Berkowitz, Birgean and Kilian (1999) for a review of this bootstrap method) and used in the construction of $\hat{\Phi}_h^{LP*}$. In constructing $\hat{\Theta}_h^{LP*}$, for each bootstrap replication, we construct A_0^{-1*} based on a draw $\widehat{\Sigma}_e^*$ from the asymptotic distribution of $\hat{\Sigma}_e^{VAR}$.[58]

A nominal $(1-\alpha)$% percentile confidence interval may be constructed, conditional on the data, as

$$P\left(\hat{\theta}_{ij,h,\alpha/2}^{LP(q)*} \leq \theta_{ij,h}^{LP(q)} \leq \hat{\theta}_{ij,h,(1-\alpha/2)}^{LP(q)*}\right) = 1-\alpha, \tag{27}$$

where $\hat{\theta}_{ij,h,\alpha/2}^{LP(q)*}$ and $\hat{\theta}_{ij,h,(1-\alpha/2)}^{LP(q)*}$ are the $\alpha/2$ and $1-\alpha/2$ quantiles of the distribution of $\hat{\theta}_{ij,h}^{LP(q)*}$. Under asymptotic normality, this interval provides a valid first-order approximation (see Efron and Tibshirani 1993). In principle, an alternative could have been to construct the symmetric percentile-t interval:

$$P\left(\hat{\theta}_{ij,h}^{LP(q)} - t_{1-\alpha}^* \frac{1}{\sqrt{T}}\widehat{\sigma}\left(\hat{\theta}_{ij,h}^{LP(q)*}\right) \leq \theta_{ij,h}^{LP(q)} \leq \hat{\theta}_{ij,h}^{LP(q)} + t_{1-\alpha}^* \frac{1}{\sqrt{T}}\widehat{\sigma}\left(\hat{\theta}_{ij,h}^{LP(q)*}\right)\right) = 1-\alpha, \tag{28}$$

where $t_{1-\alpha}^*$ denotes the $1-\alpha$ quantile of the distribution of $\left|\hat{\theta}_{ij,h}^{LP(q)*} - \hat{\theta}_{ij,h}^{LP(q)}\right| / \widehat{\sigma}\left(\hat{\theta}_{ij,h}^{LP(q)*}\right)$, and $\widehat{\sigma}\left(\hat{\theta}_{ij,h}^{LP(q)*}\right)$ is an estimate of the standard deviation of $\hat{\theta}_{ij,h}^{LP(q)*}$. The problem with this proposal is that it is not clear how to estimate the variance of the bootstrap LP impulse response estimator in constructing the studentized statistic t^*. The Newey-West estimator of Jordà (2007) cannot be used since we rely on the block bootstrap for generating bootstrap draws. Nor can the variance of the estimator be simulated by the bootstrap method, conditional on a given realization of the block-bootstrap estimator. For that reason we do not consider the percentile-t interval.

[58]Preliminary simulation experiements suggested that treating Σ_e^* as random improved the coverage accuracy of intervals compared with intervals based on the initial point estimate $\hat{\Sigma}_e^{VAR}$ for all bootstrap replications.

Joint Confidence Intervals

In addition to the pointwise confidence intervals described so far, we consider the joint confidence regions proposed by Jordà (2007). Let the $(K(H+1) \times K)$ matrix Θ be a set of impulse responses for horizon 0 to H. Let θ denote $vec(\Theta)$ and $\hat{\theta}$ the corresponding estimator. Jordà observes that

$$\sqrt{T}(\hat{\theta} - \theta) \xrightarrow{d} N(0, \Omega_\theta),$$

where Ω_θ is the $(K^2(H+1) \times K^2(H+1))$ limiting variance-covariance matrix of all structural impulse response coefficients up to some horizon H. The confidence region that contains the entire impulse response path with probability $100(1-\alpha)\%$ in repeated sampling is a multidimensional ellipsoid which cannot be displayed easily. Jordà suggests that this joint confidence region can be approximated by Scheffé's (1953) S-method. The resulting joint LP confidence interval is:

$$\left[\hat{\theta}_{ij} \pm \frac{chol(\widehat{\Omega}_{\theta,ij})}{\sqrt{T}} \sqrt{\frac{c^2_{1-\alpha}}{H+1}} i_{(H+1)} \right] \tag{29}$$

where $\theta_{ij} = [\theta_{ij,0}, \theta_{ij,1}, \theta_{ij,2}, \cdots, \theta_{ij,H}]'$, $c^2_{1-\alpha}$ is the critical value of the chi-square distribution with $H+1$ degrees of freedom, i_{H+1} is a $(H+1) \times 1$ vector of ones. Explicit expressions for $\widehat{\Omega}_{\theta,ij}$ that are applicable to both VAR and LP estimators are presented in Jordà (2007). The joint variance-covariance matrix of the impulse response coefficients is evaluated using the closed-form solution of Jordà (2007):

$$\begin{aligned} \Omega_\theta &= \left(A_0^{-1} \otimes I_{K(H+1)}\right)' \left(\left(y_t' M_x y_t\right)^{-1} \otimes \Sigma_u\right) \left(A_0^{-1} \otimes I_{K(H+1)}\right) + \\ & \quad 2\left(I_K \otimes \Phi\right))C D_K^{+} \left(\Sigma_e \otimes \Sigma_e\right) D_K^{+\prime} C' \left(I_K \otimes \Phi\right)', \end{aligned} \tag{30}$$

where Φ is the $(K(H+1)\times K)$ matrix of reduced-form impulse responses, $C = L_K'\left\{L_K\left(I_{K^2}+K_{KK}\right)\left(A_0^{-1}\otimes I_K\right)L_K'\right\}^{-1}$, and L_K is the elimination matrix, K_{KK} the commutation matrix, and D_K^+ is the duplication matrix defined in Lütkepohl (2005, Appendix 12.2). In practice, we substitute consistent estimators for the unknown expressions Φ, Σ_u, Σ_e and A_0^{-1}. We deal with possible serial correlation in the error term of the local projection by employing the Newey-West estimator of Σ_u with a truncation lag corresponding to the maximum impulse response horizon. Additional simulation evidence (not shown) suggests that somewhat tighter joint intervals may be obtained at the cost of more erratic coverage accuracy by lowering the truncation lag. Alternatively, the joint variance-covariance matrix may be evaluated using the bootstrap method conditional on the original data.

In the applications below we adopt a suggestion by Jordà and Marcellino (2009) to replace expression (29) by:

$$\left[\hat{\theta}_{ij} \pm \frac{chol(\widehat{\Omega}_{\theta,ij})}{\sqrt{T}}\sqrt{\frac{c^2_{1-\alpha}(h)}{h}}\right]_{h=0}^{H} \tag{31}$$

Additional simulation evidence shows that this modification tends to improve the finite-sample accuracy of the joint interval.

3.2.4 Evaluation Criteria

We are interested in comparing the small-sample performance of impulse responses estimated by local projections and by conventional VAR methods. One set of criteria are the effective coverage accuracy and average length of pointwise impulse responses confidence intervals. Effective coverage is defined as the relative frequency with which the confidence interval covers the true, but in practice unknown value of the impulse response in repeated sampling. Average length is the average distance between the upper and lower bounds of the confidence interval in repeated trials. In addition, we report the bias, standard deviation and mean-squared error of the impulse response point estimates to help explain differences in coverage accuracy and average length. For the joint confidence interval, we calculate the probability that the interval estimator contains the entire path of the true impulse response function in repeated sampling. We also report the average interval length. Throughout the paper, we focus on nominal 95% confidence intervals. Qualitatively similar results are obtained for nominal 68% intervals.

3.3 Simulation Evidence: Bivariate VAR(1) Model

In this section, we perform a simulation study to evaluate the relative small-sample performance of VARs and LPs in the context of a stationary VAR(1)-DGP. The results will help build intuition before we turn to a more realistic DGP in the next section. The model is:

$$y_t = \begin{pmatrix} B_{11} & 0 \\ 0.5 & 0.5 \end{pmatrix} y_{t-1} + e_t, \quad e_t \overset{iid}{\sim} N\left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 1 & 0.3 \\ 0.3 & 1 \end{pmatrix} \right) \tag{32}$$

where $B_{11} \in \{0.5, 0.9, 0.97\}$. The intercept has been normalized to zero in population. This DGP has been used widely in the literature as a benchmark (see, e.g., Griffiths and Lütkepohl 1993, Fachin and Bravetti 1996, Kilian 1998a,b; Berkowitz and Kilian 2000).

We draw 1,000 time series of length $T = 100$ from this process.[59] For each trial, we fit the VAR model and a sequence of LP models, one each for each horizon. All regression models include an intercept.[60] The lag-order for each local projection is

[59]Initially, we also considered a sample size of $T = 50$. In that case, the Σ^* draws required for the LP bootstrap method sometimes are not positive definite, making it impossible to implement the LP method. Hence, we focus on the larger sample size.

[60]With the exception of the construction of joint confidence regions in section 3.3.2 and 3.4.1, we follow Jordà's (2005) proposal of fitting individual projections rather than doing one projection for

chosen by the Akaike information criterion (AIC) with an upper bound of four lags.[61] The same criterion and upper bound is used in selecting the lag order of the fitted VAR model. For each of the impulse response estimates $\widehat{\theta}_{ij,h}^{VAR(p)}$ and $\widehat{\theta}_{ij,h}^{LP(q)}$ we construct the pointwise confidence intervals discussed earlier. All bootstrap confidence intervals are based on 2,000 bootstrap replications. For the block bootstrap method, the size of the block is set to four at all horizons. This produces the most accurate results for the LP intervals.[62] The maximum horizon H of the impulse response function is 16.

3.3.1 Pointwise Intervals

Figure 18 shows some representative results for alternative values of B_{11}. The larger B_{11}, the higher the persistence of the process. The upper panel of Figure 18 plots the effective coverage rates of nominal 95% confidence interval for $\theta_{21,h}$, where $\theta_{21,h}$ stands for the response of variable 2 to structural shock 1 in period 0 at horizon $h = 0, ..., 16$. The difference in performance between the four methods is substantial. The effective coverage rates of the standard VAR delta method interval drops quickly with increasing horizon, consistent with earlier results in the literature. At horizon 16, it is 81% for $B_{11} = 0.5$, 69% for $B_{11} = 0.9$ and 60% for $B_{11} = 0.97$. On the other hand, the effective coverage of the bias-corrected bootstrap for the VAR remains fairly close to 95% at all time horizons and for all values of B_{11}. This result illustrates that the small-sample bias in the impulse response estimates by the VAR is substantial and that the asymptotic normal approximation in the delta method interval may not

all horizons joinly. In additional sensitivity analysis we determined that the joint linear projection is less accurate than the individual linear projections in small samples.

[61]The local projections with the lag order selected for each horizon have better small sample properties than the local projections using the same lag order for all horizons. The AIC with the maximum lag order of four performs better than the AIC with the upper bound of eight lags or the SIC with either a maximum lag order of four or eight. A natural conjecture is that the performance of the LP method may be improved by enforcing greater parsimony. For example, one could set $q = 1$ in all LP regressions. Further analysis (not shown) suggests that this modification may greatly reduce the coverage accuracy of the LP interval in practice.

[62]We also investigated whether the LP bootstrap interval performed better when allowing for different block sizes by horizon, and found that a fixed block size produces more accurate intervals.

be a good approximation to the impulse response distribution in small samples.

Unlike the coverage accuracy of the asymptotic VAR interval, that of the asymptotic LP interval does not deteriorate sharply as the horizon increases, but its overall coverage accuracy declines somewhat with increasing B_{11}. Whereas the asymptotic LP interval tends to attain close to nominal coverage for $B_{11} = 0.5$, its accuracy may drop as low as 91% for $B_{11} = 0.9$ and 89% for $B_{11} = 0.97$. While far from perfect, these results are clearly superior to the VAR delta method. This finding suggests that indeed there are potential advantages to using local projections. It may be tempting to attribute these differences to the fact that the LP impulse response estimator does not require any non-linear transformations and hence is less biased. This is not the case. As Figure 19 shows, the bias of the VAR impulse response estimator is actually similar at low horizons and smaller than the bias of the LP estimator at long horizons. There is no evidence of reduced bias. Instead, the reason for the superior coverage accuracy of the LP interval is the higher variability of the LP estimator, shown in Figure 19, especially at long horizons. This difference is reflected in a substantial increase in the average interval length, as shown in the lower panel of Figure 18. This increase in variability is consistent with the less parametric nature of the LP estimator.

Returning to the first panel of Figure 18, we see that the bias-adjusted VAR bootstrap interval generally does fairly well at all horizons, with a tendency for the coverage accuracy to decline slightly as B_{11} increases. In contrast, the LP percentile bootstrap interval is even less accurate than the asymptotic LP interval in many cases, except at short horizons. The reason is that the bootstrap tends to amplify the bias in the initial point estimates, much like in the case of the Runkle (1987) VAR bootstrap method (see Kilian 1998a). To overcome that problem would require a bias adjustment of the LP regressions, but such an adjustment is not straightforward since there are no closed form solutions, nor does the block bootstrap method lend itself to bias adjustments.

To summarize, we find that, contrary to the conjecture in the introduction, there is no evidence that the LP impulse response estimator has lower bias than the VAR estimator, while the conjecture that the LP estimator has higher variance proved to be correct. Thus, local projections tend to deliver less accurate point estimates of impulse responses in terms of the MSE. As for inference, we found that the bias-corrected bootstrap interval for VAR models dominates all other pointwise confidence intervals in this simulation example. Moreover, there is no evidence that bootstrapping improves the accuracy of LP confidence intervals.

3.3.2 Joint Intervals

Table 17 investigates the coverage accuracy and average length of the corresponding asymptotic joint intervals. In computing the coverage accuracy of the nominal 95% joint interval, we compute the relative frequency with which the interval estimator includes the entire true response function within the interval bounds. Thus, there is one coverage rate for each of the four impulse response functions. We find that the joint VAR interval is typically more accurate than the joint LP interval. The coverage rate of the nominal 95% joint VAR interval ranges from 86% to 98%. The corresponding results for the joint LP interval range from 80% to 91%. At the same time, the average interval length of the joint LP interval is always greater than that of the joint VAR interval, often by a factor of more than two. Thus, neither type of joint interval necessarily comes close to nominal coverage, even in this simple model.

These estimates are based on the closed-form solution of the asymptotic variance provided in Jordà (2007). We also investigated the extent to which the accuracy of these intervals can be improved by estimating the joint variance-covariance matrix of the impulse response coefficient estimators by bootstrap methods. We found that substituting suitably constructed bootstrap estimates of this matrix based on 2,000 bootstrap draws did not improve the accuracy of either the VAR or LP interval. The relative ranking of the two methods remained unchanged. These additional results are not shown to conserve space.[63]

3.3.3 Larger Sample Sizes

The preceding analysis highlighted some practical limitations of the LP method in small samples. It is clear that, as the sample size increases, the performance of the LP point and interval estimators will improve. As expected, doubling the sample size from $T = 100$ to $T = 200$ greatly reduces the bias and variance of the LP estimator and improves the accuracy of the pointwise asymptotic LP interval to near nominal coverage even for $B_{11} = 0.97$, while reducing its average length. Although the asymptotic LP intervals for $T = 200$ are about as accurate as the VAR bootstrap intervals, they remain about three times as wide on average, however. Thus, one would still prefer the VAR-based interval for pointwise inference. For the joint intervals coverage improves for $T = 200$, but not dramatically so. The accuracy of the LP and VAR intervals becomes more similar. The joint VAR interval has coverage rates between 88% and 98%; the joint LP interval between 88% and 96%. For a given response function, the ranking by accuracy is in general ambiguous, but the joint LP interval always is considerably wider on average.

3.4 Four-Variable VAR(12) Model

The preceding results were confined to a stylized VAR(1) model. In this section we demonstrate that our main conclusions continue to hold in a realistic example with many lags and variables. Our example is a prototypical partially identified four-

[63] We did not attempt to construct joint intervals based on the bootstrap. While it would not be difficult to simulate the joint distribution of the pointwise impulse response estimators, the resulting confidence region could not be displayed easily. It was this fact that prompted Jordà (2007) to propose the use of Scheffé's (1953) S-method resulting in the asymptotic joint interval defined in section 3.2.

variable VAR model of the type commonly employed in the analysis of monetary policy shocks (see, e.g., Christiano, Eichenbaum and Evans 1999). We postulate a VAR(12) model with intercept for $y_t = [gap_t,\ \pi_t,\ \pi_t^{RPCOM},\ i_t]'$. Underlying this model is the notion that the Federal Reserve sets the interest rate (i_t), conditional on all past data, as a function of the current inflation rate (π_t) and output gap (gap_t). We follow the literature in augmenting the model with the growth rate in real industrial commodity prices, as a leading indicator of inflationary pressures (π_t^{RPCOM}). The presumption is that this additional variable helps alleviate the well-known price puzzle. The model is semi-structural in that only the monetary policy shock is identified. As is standard in this literature, the identifying assumption is that there is no contemporaneous feedback from policy decisions to the output gap, to commodity prices, or to the inflation rate. We specify the model at monthly frequency since the identifying assumptions are more credible at monthly than at quarterly frequency.

This system is similar to models discussed in Christiano et al. (1999) and closely resembles the VAR model of Bernanke and Gertler (1995). One difference between their models and this model is that our measure of output is broader and clearly stationary and that in our model the price level is specified in log-differences. This transformation ensures that the model is stationary, if still persistent. That fact is important since the maintained assumption in this paper is stationarity. The sample period is January 1970 through December 2007. We also repeated our analysis for subsamples. The simulation results were qualitatively similar and are not reported.

Our measure of inflation is based on the seasonally adjusted monthly CPI for all urban consumers. Real commodity price inflation is constructed as the change in the Commodity Research Board's price index for raw industrials adjusted for CPI inflation. The Federal Funds rates serves as our proxy for the interest rate. Our measure of the real output gap is the CFNAI, a weighted average of a multitude of monthly indicators of U.S. real economic activity, constructed by the Federal Reserve Bank of

Chicago. This is a principal components index based on 85 real indicators including measures of production, income, employment, and consumption. It is constructed to have an average value of zero and a standard deviation of one and is stationary by construction. The CFNAI can be interpreted as a measure of the U.S. business cycle. The use of the CFNAI has several advantages relative to other output measures. First, real GDP data are not available at monthly frequency and industrial production data capture only a small and declining share of output. Second, it is well established that the Federal Reserve considers many measures of real output rather than one time series only (see, e.g., Evans 1999). The use of principal components allows us to capture a broader set of business cycle indicators. Third, conventional measures of output tend to imply implausibly large and persistent effects of monetary policy shocks on output, whereas our measure generates the expected temporary response. This result is in line with recent work stressing the importance of incorporating information from larger data sets in VAR models (see, e.g., Stock and Watson 2005; Banbura, Giannone and Reichlin 2008).

3.4.1 Simulation Evidence

Figure 20 shows effective coverage rates and average lengths of alternative pointwise confidence intervals up to a horizon of $H = 24$. These are obtained by generating 1,000 trials of the same length as the original data from the VAR(12) model fitted on the actual data. In that simulation exercise, the VAR error term is assumed to be Gaussian. The lag orders of the fitted regression models are obtained using the AIC with an upper bound of twelve lags.

We focus on the responses of the output gap, of CPI inflation and of real commodity price inflation to an unanticipated monetary policy tightening. Figure 20 reveals severe coverage deficiencies for the LP bootstrap interval especially at long horizons. Its coverage rates may drop as low as 75%. The asymptotic LP method is fairly accurate at short horizons (except on impact), but its coverage accuracy also deteriorates at longer horizons. For example, for the output gap its coverage rate may drop as low as 90%, for the inflation rate as low as 88% and for real commodity price inflation as low as 82%. In contrast, both the asymptotic VAR interval and the bias-adjusted bootstrap VAR interval are consistently quite accurate. If anything, their coverage is excessive. Moreover, there is little to choose between the two VAR intervals in terms of their average length. The VAR intervals not only tend to be more accurate, but also systematically shorter than the LP intervals. The asymptotic LP interval is sometimes more than twice as wide on average as the other intervals. Even the bootstrap LP interval, however, tends to be wider than the VAR intervals.

Similarly, when it comes to point estimates of the impulse responses, Figure 21 shows that the LP estimator has higher bias in most cases as well as higher variance, resulting in unambiguously higher MSEs. These results are very much consistent with the insights obtained from the stylized VAR(1) model.

Table 18 shows the corresponding results for the joint interval. Here the results differ somewhat from the VAR(1) data generating process in that the joint LP interval has typically more accurate coverage with rates between 88% and 94%. In contrast, the accuracy of the joint VAR interval is typically lower and more erratic with rates between 61% and 88%. As in the VAR(1) example, the joint LP interval is much wider on average often by a factor of more than three. This conclusion is further supported by results for subsamples, which generally produced much less accurate results for both methods. For example, if we restrict the sample to the pre-Greenspan period of 1970.1-1987.8, joint coverage rates drop as low as 33% for the VAR interval and as low as 36% for the LP interval. Moreover, the rankings of the two methods are mixed. We infer that joint intervals are unlikely to be reliable in practice, and that neither joint interval is clearly preferred over the other.

3.4.2 Empirical Application

We conclude this section with an illustration of how different methods of inference and choices of models may affect the consensus about the responses of macroeconomic aggregates to an unanticipated monetary policy tightening. The data and model are identical with the model used as a data generating process in this section. The first row of Figure 22 shows the responses estimated by the VAR model and the second row the responses estimated by local projections. The point estimates are plotted together with conventional pointwise nominal 95% confidence intervals. The VAR estimates suggest that the output gap temporarily turns negative with a full recovery within two years. The temporary decline is highly statistically significant. Real commodity price inflation also drops significantly. There is no evidence that CPI inflation is significantly reduced. Rather the inflation response exhibits the well known price puzzle with a statistically significant peak after two months. There is little to choose between bias-corrected bootstrap and delta method intervals in this application.

The LP estimates in the second row paint a very similar picture. The main difference is that the intervals tend to be substantially wider, lowering the statistical significance, and that the estimates are more erratic. Both findings are expected based on the earlier Monte Carlo simulation evidence. The pointwise bootstrap LP intervals are so wide at short horizons that not even the price puzzle remains significant. In contrast, based on the asymptotic LP interval the price puzzle remains.

Figure 23 compares the pointwise and joint confidence intervals for the same example. The first row illustrates that the joint VAR intervals are similar to pointwise intervals at short horizons, but systematically wider at long horizons. A similar pattern applies to the LP intervals in the second row. At short horizons the joint and pointwise asymptotic intervals are quite similar. At longer horizons, discrepancies emerge, as the joint interval widens disproportionately. Compared with the joint VAR intervals, these intervals are substantially wider, as predicted by the simulation study.

These differences do not necessarily affect the substantive conclusions, however, because statistically significant point estimates tend to be concentrated at short horizons, when joint and pointwise intervals tend to be similar. For the VAR results, for none of the three response functions the use of joint intervals overturns the finding of a significant response function. For the LP results, both the price puzzle and the output contraction remain statistically significant, if barely so, even using the joint interval. Thus, at least in this empirical example, the use of LP models as opposed to VAR models and the use of joint intervals as opposed to pointwise intervals makes little difference for the main conclusions.

3.5 Approximate VAR Models

Based on the simulation results presented so far, there is no compelling reason to abandon traditional VAR methods of constructing impulse response estimates in favor of the LP method. As long as the finite lag order VAR model provides a good approximation to the stationary data generating process, the LP estimator suffers from greater small-sample bias and higher variance, resulting in very wide, yet often still inaccurate confidence intervals and erratic point estimates. An interesting avenue for future research is to investigate how poor the vector autoregressive approximation has to be for the LP method to become an attractive alternative to VAR approximations. For example, more general linear stationary processes can be represented as a VAR(∞) model and approximated by a sequence of finite-lag order vector autoregressions (see, e.g., Lütkepohl and Poskitt 1991; Inoue and Kilian 2002).

In this section, we provide some preliminary evidence based on a tri-variate invertible VARMA(1,1)-DGP used as an example in Braun and Mittnik (1993) and Inoue and Kilian (2002). The model includes quarterly investment growth, deflator inflation and the commercial paper rate in this order:

$$y_t = A_1 y_{t-1} + \varepsilon_t + M_1 \varepsilon_{t-1},$$

where $A_1 = \begin{bmatrix} 0.5417 & -0.1971 & -0.9395 \\ 0.0400 & 0.9677 & 0.0323 \\ -0.0015 & 0.0829 & 0.8080 \end{bmatrix}$, $M_1 = \begin{bmatrix} -0.1428 & -1.5133 & -0.7053 \\ -0.0202 & 0.0309 & 0.1561 \\ 0.0227 & 0.1178 & -0.0153 \end{bmatrix}$,

and $\varepsilon_t \sim NID(0, PP')$ with $P = \begin{bmatrix} 9.2352 & 0 & 0 \\ -1.4343 & 3.6070 & 0 \\ -0.7756 & 1.2296 & 2.7555 \end{bmatrix}$.

This model can be represented as a recursively identified VAR(∞) process which can be approximated using either local projections or a finite-lag order VAR model. We follow Inoue and Kilian (2002) in studying the response of the model variables to an innovation in the commercial paper rate. We focus on the asymptotic LP interval and the bias-adjusted bootstrap algorithm for VAR models.[64] The sample size is $T = 200$. Figure 24 shows results for approximating lag orders of $p = 5$ and $q = 5$. The simulation results are remarkably robust to changes in these lag orders. Only for very low lag orders, the coverage accuracy deteriorates.

The bias-adjusted bootstrap method performs quite well, as illustrated in Figure 24. Based on the approximating VAR(5) model, its coverage rate is near the nominal coverage for all three responses. Interestingly, Figure 24 suggests that, for a suitably large choice of q, the asymptotic LP interval is just as accurate, but its average width is systematically wider by a factor of about three. Thus, the VARMA-DGP results are quite similar to the results we obtained earlier for finite-lag order VAR models for large T. This tentative evidence suggests that it is not clear that there are advantages to the LP approach, even if the VAR model is merely an approximation to the data generating process.

[64] Inoue and Kilian (2002) discuss the validity of the bootstrap for stationary VAR(∞) processes and show that this specific bootstrap approach performs better than the delta method interval proposed in Lütkepohl and Poskitt (1991).

3.6 Conclusion

Local projections methods are a promising recent development in the literature on impulse response analysis, but little is known about their finite sample performance and their merits relative to more conventional VAR-based methods. In this paper, we compared the small-sample performance of impulse response confidence intervals and point estimates based on local (linear) projections and VAR models. We explored and compared alternative approaches to implementing the LP method in stationary ennvironments and developed suitable bootstrap methods of inference.

Our main objective was to investigate the conjecture that LP intervals may help resolve the long-standing problem of bias driven by the nonlinearity of the VAR impulse-response estimator. This bias tends to undermine Gaussian approximations to the finite-sample distribution of the impulse response in vector autoregressions, resulting in confidence intervals with poor coverage accuracy. Since LP impulse responses can be represented as slope coefficients in a linear model, our conjecture was that confidence intervals based on the LP model might be more accurate in practice than VAR based intervals.

We showed that this conjecture is not correct. In particular, we found that the bias of the LP estimator is greater than the bias of the VAR impulse response estimator, notwithstanding the linearity of the LP estimator in the slope parameters. Combined with its excessive variability, the LP point estimator proved to be very unreliable. The accuracy of the LP interval estimator tended to be erratic in small samples. Although the asymptotic LP interval was more accurate than some VAR-based alternatives such as the delta method interval, neither the asymptotic nor the bootstrap LP interval proved more accurate than bias-adjusted bootstrap intervals for VAR models. We concluded that there is no compelling reason to abandon traditional VAR methods of constructing impulse response estimates in favor of the LP method, as long as the finite lag order VAR model provides a good approximation to the data generating process. While the accuracy of LP estimators quickly improves with increasing sample size, the average width of the asymptotic LP interval far exceeds that of bias-adjusted VAR bootstrap intervals with similar accuracy. Thus, even for large samples there are no apparent advantages to the LP method. This result tends to hold even when the data are generated by a stationary VAR(∞) process.

We also investigated the reliability of the joint impulse response confidence intervals proposed by Jordà (2007) for both VAR and LP models. Such intervals are of great potential interest for applied work, given the dependence of impulse response estimates across horizons. We concluded that these joint intervals were not accurate enough to be recommended for realistic applications. This conclusion held whether these intervals were constructed based on the VAR or the LP approach. An empirical application illustrated that joint VAR intervals tend to be similar to pointwise intervals at short horizons. Only at longer horizons they become substantially wider than pointwise intervals. Since most statistically significant results in applied work are obtained at short horizons, there is reason to believe that in most cases the use of joint intervals is unlikely to overturn the substantive conclusions of studies based on more conventional pointwise intervals.

References

Andrews, D.W.K., and H.-Y. Chen (1994), "Approximately Median-Unbiased Estimation of Autoregressive Models," *Journal of Business and Economic Statistics*, 18, 139-165.

Banbura, M., D. Giannone and L. Reichlin (2008), "Large Bayesian VARs," Working Paper No. 966, European Central Bank.

Benkwitz, A., H. Lütkepohl , and M.H. Neumann. (2000), "Problems Related to Confidence Intervals for Impulse Responses of Autoregressive Processes," *Econometric Reviews*, 19, 69-103.

Berkowitz, J., I. Birgean, and L. Kilian (1999), "On the Finite-Sample Accuracy of Nonparametric Resampling Algorithms for Economic Time Series," in T.B. Fomby and C. Hill (eds.): *Advances in Econometrics: Applying Kernel and Nonparametric Estimation to Economic Topics*, vol. 14, 1999, JAI Press, Connecticut, 77-107.

Berkowitz, J., and L. Kilian (2000), "Recent Developments in Bootstrapping Time Series," *Econometric Reviews*, 19, 1-48.

Bernanke, B.S. and M. Gertler (1995), "Inside the Black Box: The Credit Channel of Monetary Policy Transmission," *Journal of Economic Perspectives*, 9, 27-48.

Braun, P. A., and S. Mittnik (1993), "Misspecifications in Vector Autoregressions and Their Effects on Impulse Responses and Variance Decompositions," *Journal of Econometrics*, 59, 319-41.

Christiano, L.J., M. Eichenbaum and C.L. Evans (1999), "Monetary Policy Shocks: What Have We Learned and to What End?" in: *Handbook of Macroeconomics,* 1, Edited by J.B. Taylor and M. Woodford.

Davidson, R. (2000), "Comment on 'Recent Developments in Bootstrapping Time Series'," *Econometric Reviews*, 19, 49-54.

Evans, C.L. (1999), "If You Were the Central Banker, How Many Data Series Would You Watch? An empirical analysis," unpublished manuscript, Federal Reserve Bank of Chicago.

Fachin, S. and L. Bravetti (1996), "Asymptotic Normal and Bootstrap Inference in Structural VAR Analysis," *Journal of Forecasting*, 15, 329-341.

Forni, M. and L. Reichlin (1998), "Let's Get Real: A Factor Analytical Approach to Disaggregated Business Cycle Dynamics," *Review of Economic Studies*, 65, 453-473.

Goncalves, S., and L. Kilian (2004), "Bootstrapping Autoregressions with Conditional Heteroskedasticity of Unknown Form," *Journal of Econometrics*, 123, 89-120.

Goncalves, S., and L. Kilian (2007), "Asymptotic and Bootstrap Inference for AR(∞) Processes with Conditional Heteroskedasticity," *Econometric Reviews*, 26, 609-641.

Gospodinov, N. (2004), "Asymptotic Confidence Intervals for Impulse Responses of Near-integrated Processes", *Econometrics Journal*, 7, 505-527.

Griffiths, W., and H. Lütkepohl (1993), "Confidence Intervals for Impulse Responses from VAR Models: A Comparison of Asymptotic Theory and Simulation Approaches," manuscript, Institut für Statistik und Ökonometrie, Humboldt-Universität, Berlin.

Hansen, B. (1999), "Bootstrapping the Autoregressive Model", *Review of Economics and Statistics*, 81, 594-607.

Inoue, A., and L. Kilian (2002a), "Bootstrapping Autoregressive Processes with Possible Unit Roots," *Econometrica*, 70, 377-391.

Inoue, A., and L. Kilian (2002b), "Bootstrapping Smooth Functions of Slope Parameters and Innovation Variances in VAR(∞) Models," *International Economic Review*, 43, 309-332.

Inoue, A., and L. Kilian (2003), "The Continuity of the Limit Distribution in the Parameter of Interest is not Essential for the Validity of the Bootstrap," *Econometric Theory*, 19, 944-961.

Jordà, Ò. (2005), "Estimation and Inference of Impulse Responses by Local Projections," *American Economic Review*, 95, 161-182.

Jordà, Ò. (2007), "Simultaneous Confidence Regions for Impulse Responses," *Review of Economics and Statistics*, forthcoming.

Jordà, Ò., and M. Marcellino (2009), "Path Forecast Evaluation," *Journal of Applied Econometrics*, forthcoming.

Kilian, L. (1998a), "Small-sample confidence intervals for impulse response functions," *Review of Economics and Statistics*, 80, 218–230.

Kilian, L. (1998b), "Confidence intervals for impulse responses under departures from normality," *Econometric Reviews*, 17, 1–29.

Kilian, L. (1998c), "Accounting for Lag Order Uncertainty in Autoregressions: The Endogenous Lag Order Bootstrap Algorithm," *Journal of Time Series Analysis*, 19, 531-548.

Kilian, L. (1999), "Finite-sample properties of percentile and percentile-t bootstrap confidence intervals for impulse responses," *Review of Economics and Statistics*, 81, 652–660.

Kilian, L. (2001), "Impulse Response Analysis in Vector Autoregressions with Unknown Lag Order," *Journal of Forecasting*, 20, 161-179.

Kilian, L., and P.-L. Chang (2000), "How Accurate are Confidence Intervals for Impulse Responses in Large VAR Models?" *Economics Letters*, 69, 299-307.

Lütkepohl, H. (1990), "Asymptotic distributions of impulse response functions and forecast error variance decompositions of vector autoregressive models," *Review of Economics and Statistics*, 72, 116–125.

Lütkepohl, H., and D.S. Poskitt (1991), "Estimating orthogonal impulse responses via vector autoregressive models," *Econometric Theory*, 7, 487–496.

Pope, A.L. (1990), "Biases of Estimators in Multivariate Non-Gaussian Autoregressions," *Journal of Time Series Analysis*, 1, 249-258.

Pesavento, E. and B. Rossi (2006), "Small-Sample Confidence Intervals for Multivariate Impulse Response Functions at Long Horizons," *Journal of Applied Econometrics*, 21, 1135-1155.

Pesavento, E. and B. Rossi (2007), "Impulse Response Confidence Intervals for Persistent Data: What Have We Learned?" *Journal of Economic Dynamics and Control*, 31, 2398-2412.

Rudebusch, G.D. (1992), "Trends and Random Walks in Macroeconomic Time Series: A Re-Examination," *International Economic Review*, 33, 661-680.

Runkle, D.E. (1987), "Vector Autoregression and Reality," *Journal of Business and Economic Statistics*, 5, 437–442.

Scheffé, H. (1953), "A Method for Judging All Contrasts in the Analysis of Variance," *Biometrika*, 40, 87-104.

Stine, R.A.(1987), "Estimating Properties of Autoregressive Forecasts," *Journal of the American Statistical Association*, 82, 1072-1078.

Stock, J. H. (1991), "Confidence Intervals for the Largest Autoregressive Root in U.S. Economic Time Series," *Journal of Monetary Economics*, 28, 435-460.

Stock, J.H. and M.W. Watson (2005), "Implications of Dynamic Factor Models for VAR Analysis," mimeo, Department of Economics, Harvard University.

Wright, J.H.(2000), "Confidence Intervals for Univariate Impulse Responses with a Near Unit Root", *Journal of Business and Economic Statistics*, 18, 368-373.

Table 16: Models for Estimating Impulse Responses and Methods of Inference

	Inference		
	Pointwise Interval		Joint Interval
Model	Asymptotic	Bootstrap	Asymptotic
VAR	(1)	(2)	(5)
LP	(3)	(4)	(6)

Figure 18: Coverage Rates and Average Lengths of 95% Pointwise Confidence Intervals for $\theta_{21,h}$

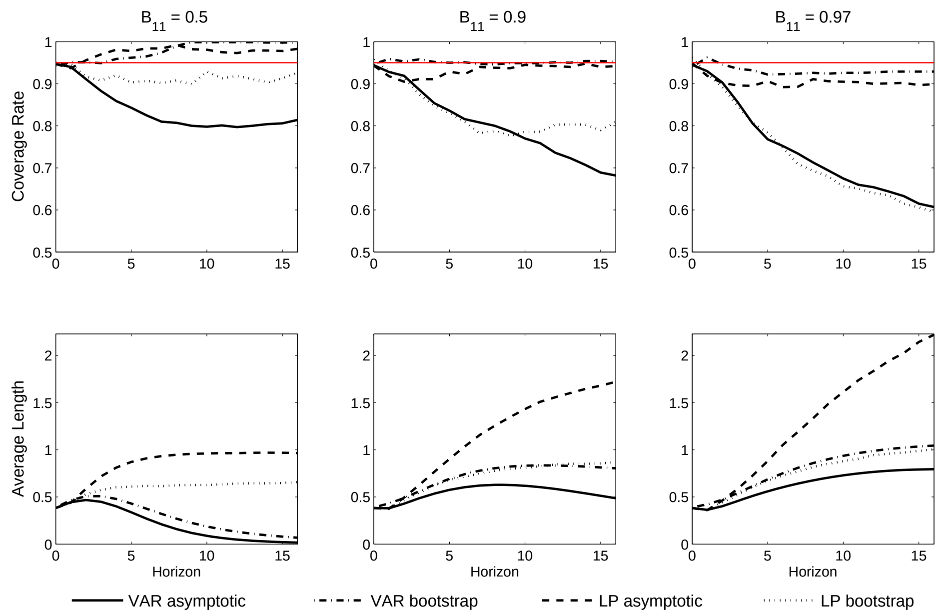

Notes: Simulation results based on 1,000 trials of length T=100 from the VAR(1) DGP described in text. VAR asymptotic denotes the asymptotic delta method for VAR impulse responses. VAR bootstrap refers to the bias-corrected bootstrap method for VAR impulse responses. LP asymptotic is the asymptotic interval for LPs. LP bootstrap refers to a block bootstrap interval for LPs. All lag orders are selected by the AIC with an upper bound of four lags for all methods.

Figure 19: Bias, Standard Deviation, and MSE of $\hat{\theta}_{21,h}$

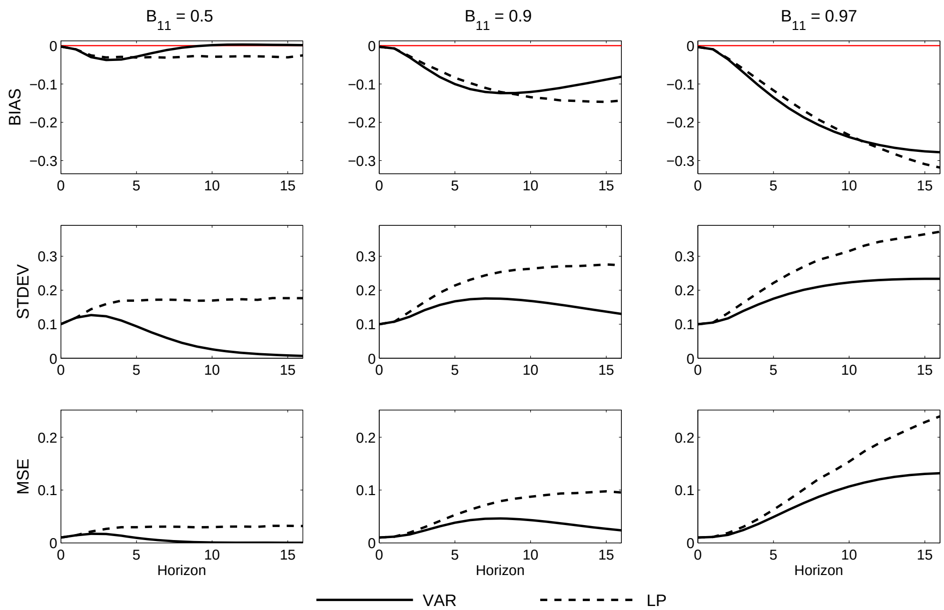

Notes: Based on 5,000 trials.

Table 17: Coverage Rates and Average Lengths of Asymptotic 95% Joint Confidence Intervals for θ_{ij} in the VAR(1) DGP

	VAR			LP		
	B_{11}=0.5	B_{11}=0.9	B_{11}=0.97	B_{11}=0.5	B_{11}=0.9	B_{11}=0.97
			Coverage Rate			
θ_{11}	0.868	0.882	0.859	0.816	0.874	0.849
θ_{12}	0.977	0.979	0.980	0.822	0.905	0.904
θ_{21}	0.899	0.907	0.896	0.867	0.872	0.843
θ_{22}	0.892	0.897	0.884	0.798	0.800	0.798
			Average Length			
θ_{11}	0.513	1.395	1.829	1.239	3.323	4.195
θ_{12}	0.494	1.351	1.767	1.203	3.202	4.025
θ_{21}	0.782	1.428	1.751	1.886	3.310	3.898
θ_{22}	0.764	1.369	1.664	1.853	3.249	3.809

Notes: Simulation results based on 1,000 trials from the VAR(1) DGP described in text.

Figure 20: Coverage Rates and Average Lengths of 95% Pointwise Confidence Intervals for Responses to a Monetary Tightening

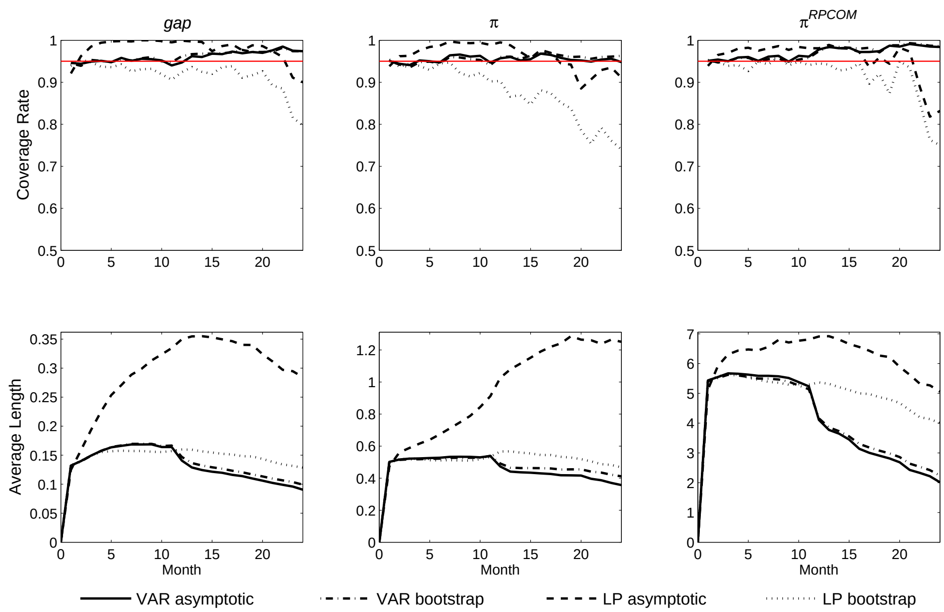

Notes: Simulation results are based on 1,000 trials of length 456 from the VAR(12) model described in the text. Gap denotes the CFNAI, π is CPI inflation, and π^{RPCOM} is real commodity price inflation. *VAR asymptotic* refers to the delta method interval for VAR impulse responses. *VAR bootstrap* refers to the bias-adjusted bootstrap method for VAR models. *LP asymptotic* refers to asymptotic interval for LP models. *LP bootstrap* refers to the block bootstrap interval for LP models. Lag orders are selected by the AIC with an upper bound of 12 lags in all cases. Since there is no uncertainty about the impact response of these variables, we do not construct a coverage rate for horizon 0.

Table 18: Coverage Rates and Average Lengths of Asymptotic 95% Joint Confidence Intervals for Responses to a Monetary Tightening in the VAR(12)-DGP

	VAR	LP
	Coverage Rate	
gap	0.877	0.920
π	0.764	0.914
π^{RPCOM}	0.614	0.936
i	0.874	0.881
	Average Length	
gap	0.265	0.853
π	1.020	2.981
π^{RPCOM}	5.082	16.589
i	0.832	2.403

Figure 21: Bias, Standard Deviation, and MSE of Impulse Responses to a Monetary Tightening

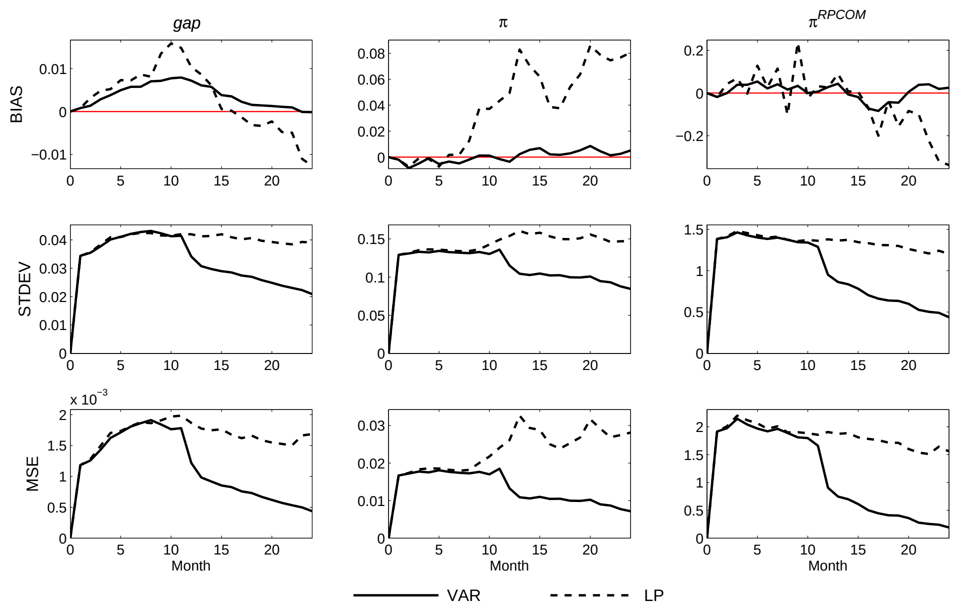

Notes: Based on 5,000 trials.

Figure 22: Responses to a Monetary Tightening with 95% Pointwise Confidence Intervals

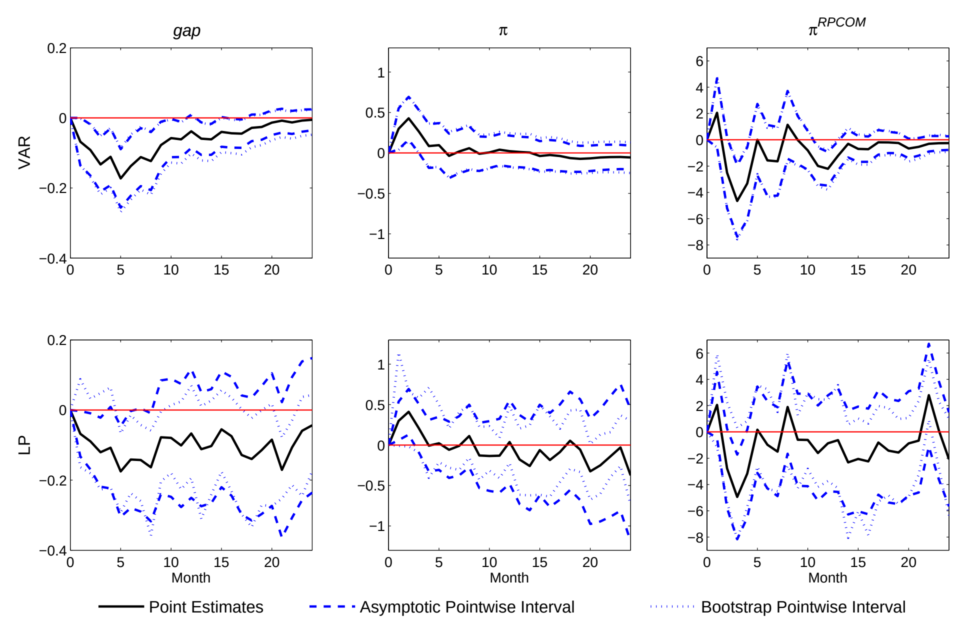

Notes: The lag orders were chosen by the AIC with an upper bound of 12 lags.

Figure 23: Responses to a Monetary Tightening with 95% Pointwise and Joint Confidence Intervals

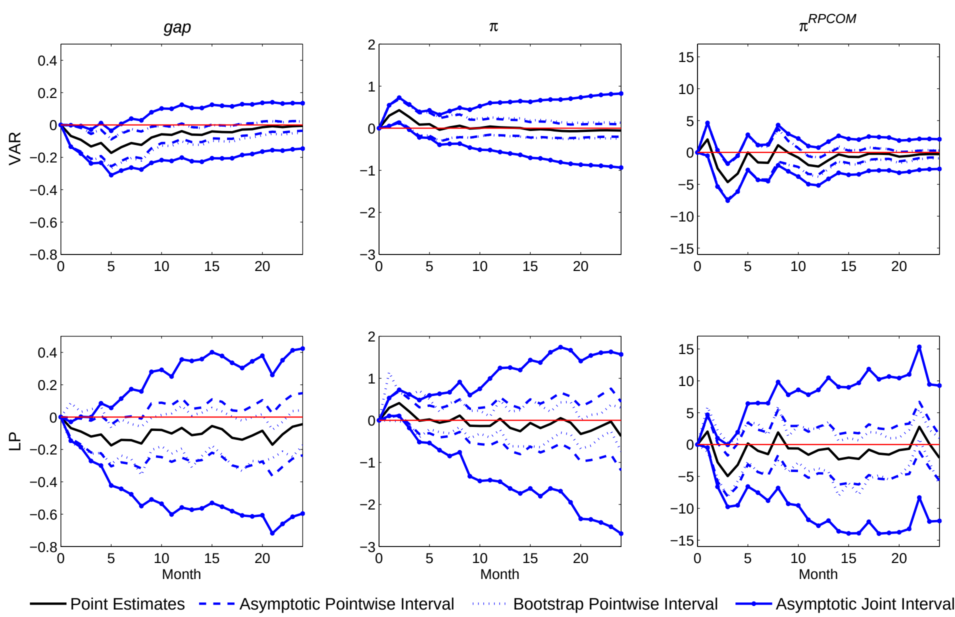

Notes: See Figure 22.

Figure 24: Coverage Rates and Average Lengths of 95% Pointwise Confidence Intervals for Responses to an Interest Rate Innovation

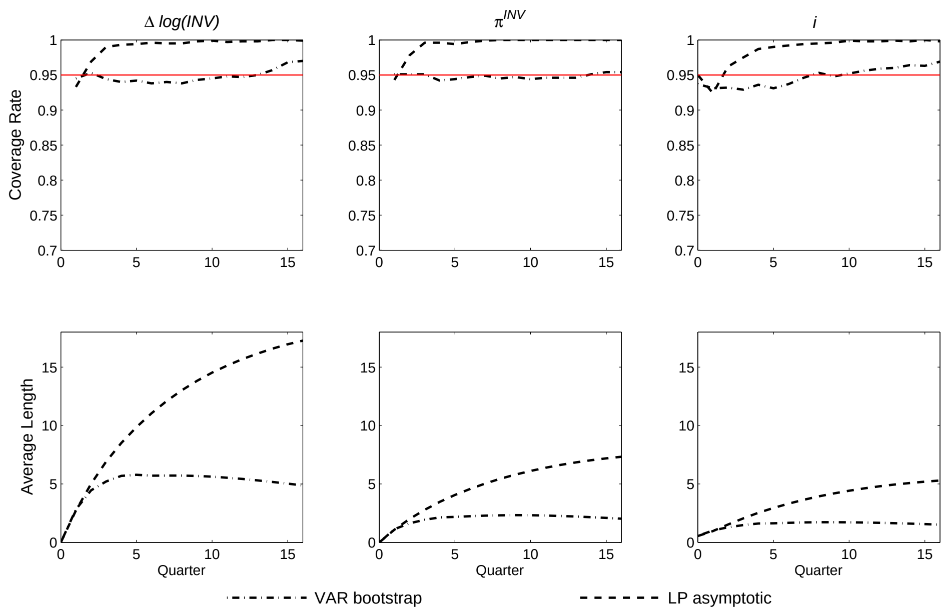

Notes: Simulation results for $T = 200$ based on 1,000 trials from the VARMA(1,1)-DGP described in Inoue and Kilian (2002). *VAR bootstrap* refers to the bias-adjusted bootstrap method for VAR models. *LP asymptotic* refers to asymptotic interval for LP models.The approximating lag orders are $p = 5$ and $q = 5$. The results are robust to reasonable changes in the lag order, as long as the lag orders are not too small.